ISSCC 2016

SESSION 4
Digital Processors

14nm 6th Generation Core Processor SOC with Low-Power Consumption and Improved Performance

Eyal Fayneh, Marcelo Yuffe, Ernest Knoll,
Michael Zelikson, Muhammad Abozaed,
Yair Talker, Ziv Shmuely, Saher Abu Rahme
Intel Corporation Haifa, Israel

Content

- Chip overview
- General power optimization technologies
  - Power delivery network
  - Clocking
  - I/O
- Power technologies – Scenario specific optimization
  - Video playback
- Power performance Si results
- Summary

Chip Overview

- SKL is a family of products manufactured using an Intel 14nm tri-gate CMOS, 12 metal layer (above M0) technology
- 2 or 4 IA cores
- System Agent (SA) including all Memory Subsystem & other IP's
- Shared last level cache (LLC, 1MByte/core)
- Scalable graphic processor (GP) with 24, 48 or 72 execution units
- Image processing unit (IPU, supporting 4 cameras simultaneously)

Chip Overview

- 2 Memory channels supports DDR3/LPDDR3/DDR4
- Display engine (DE) and 3 display I/O ports configurable to eDP, DP or HDMI
- In mobile SKUs the peripheral control hub (PCH) resides in the same package (MCP) with the CPU and communicate through an On Package IO (OPIO) bus
- In Desk-Top (DT) the PCH resides on the platform

Die floor plan & block diagram (2+2)

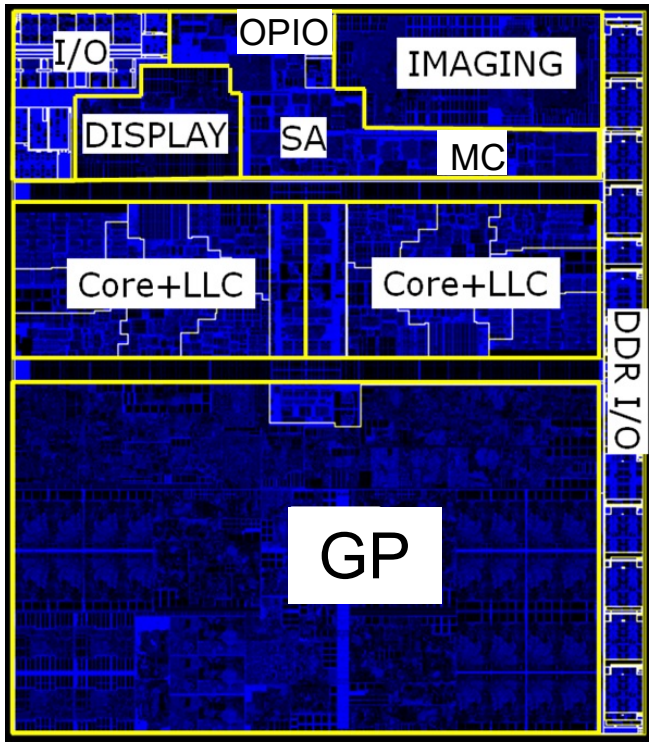

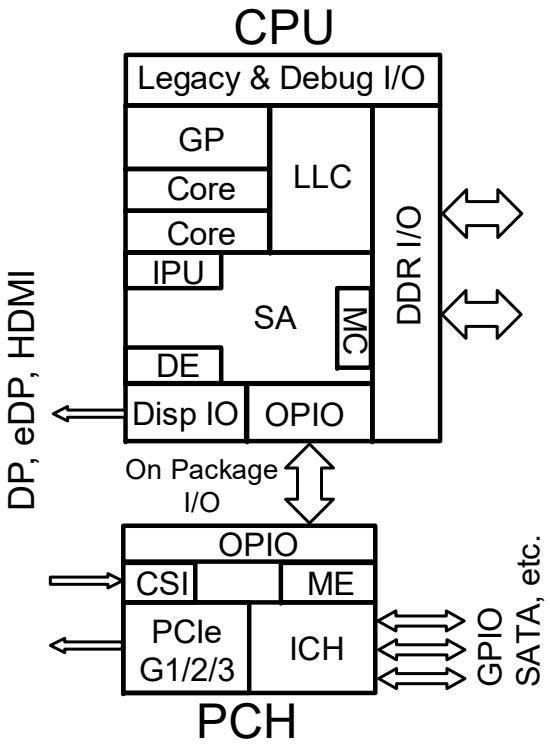

Die floor plan & block diagram (2+2)

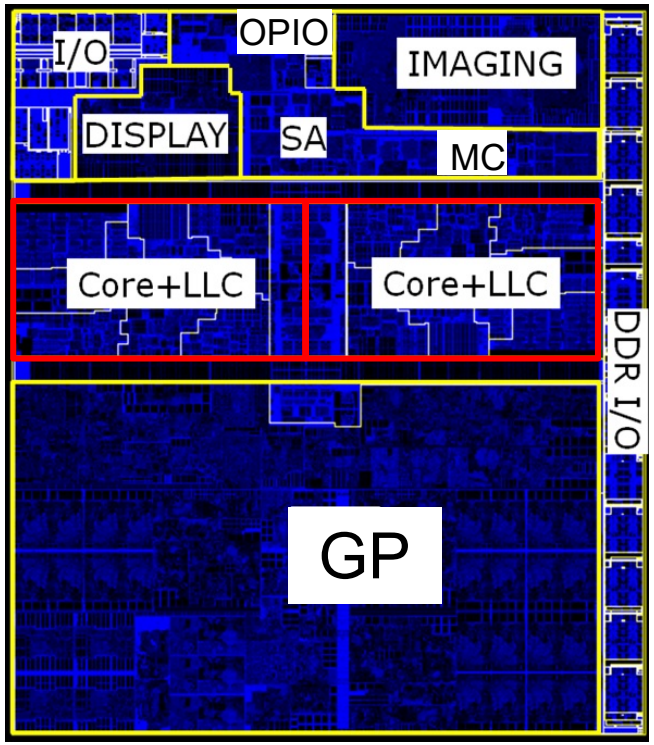

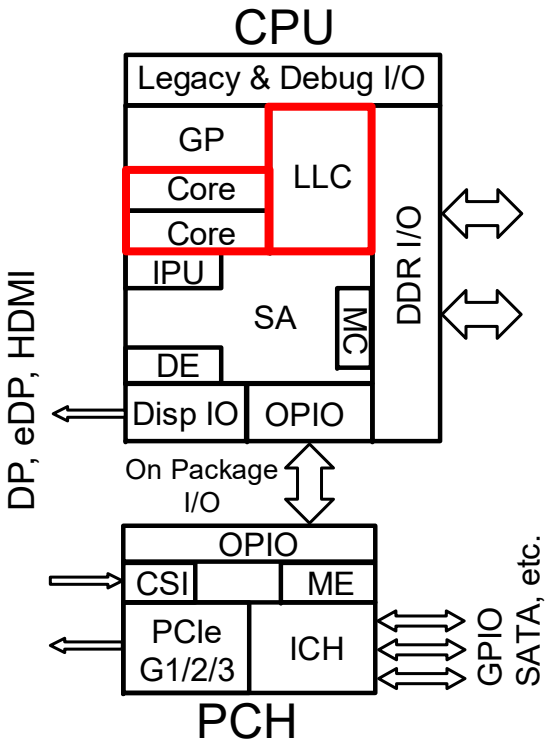

Die floor plan & block diagram (2+2)

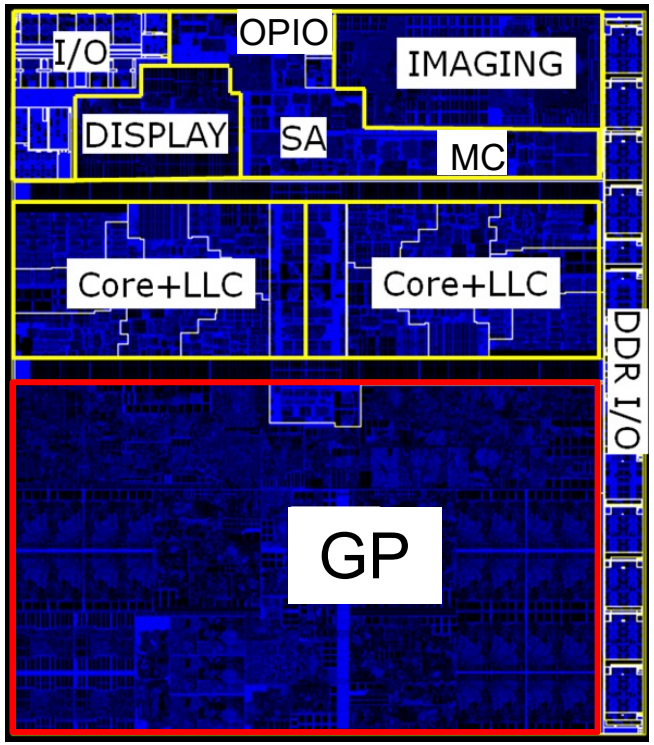

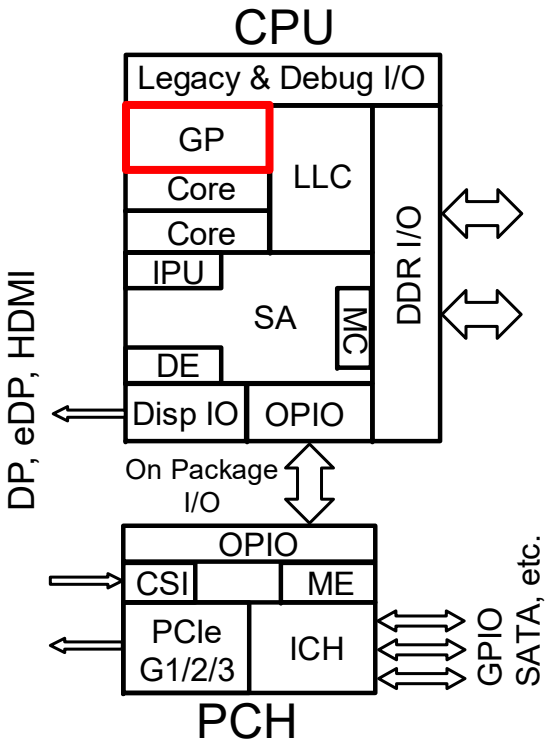

Die floor plan & block diagram (2+2)

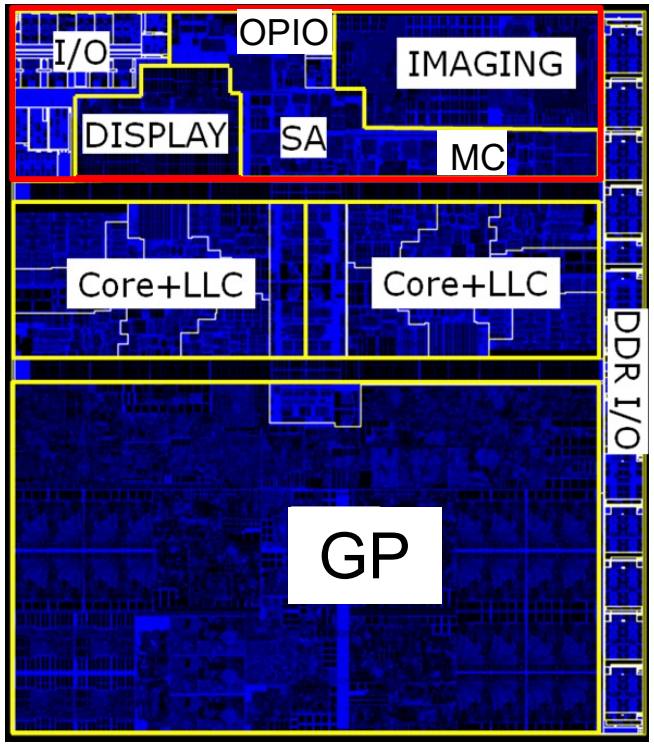

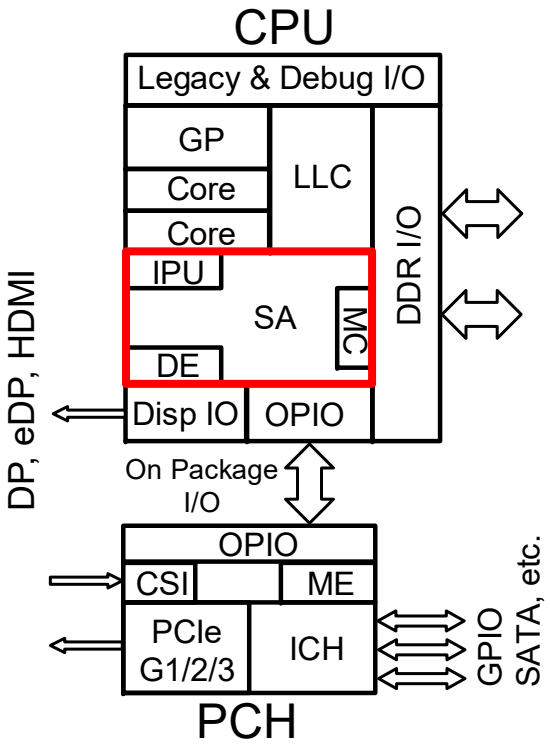

Die floor plan & block diagram (2+2)

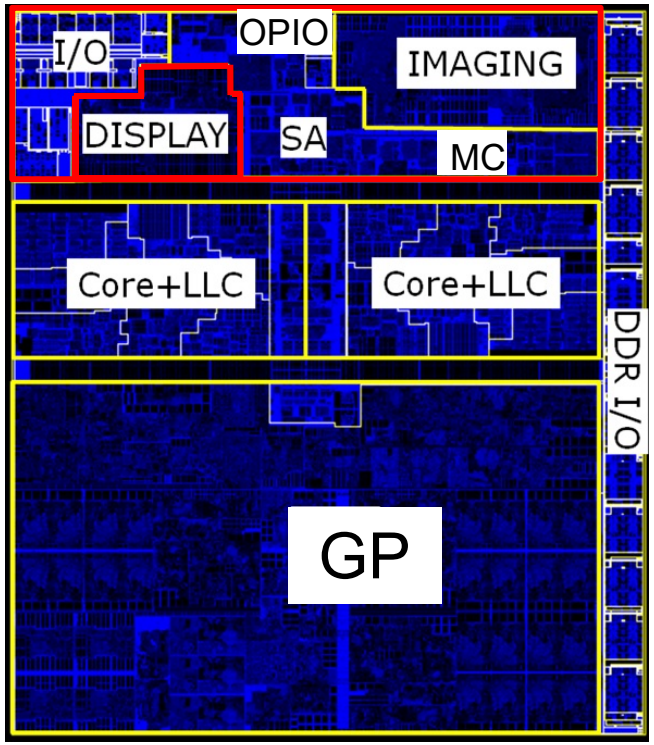

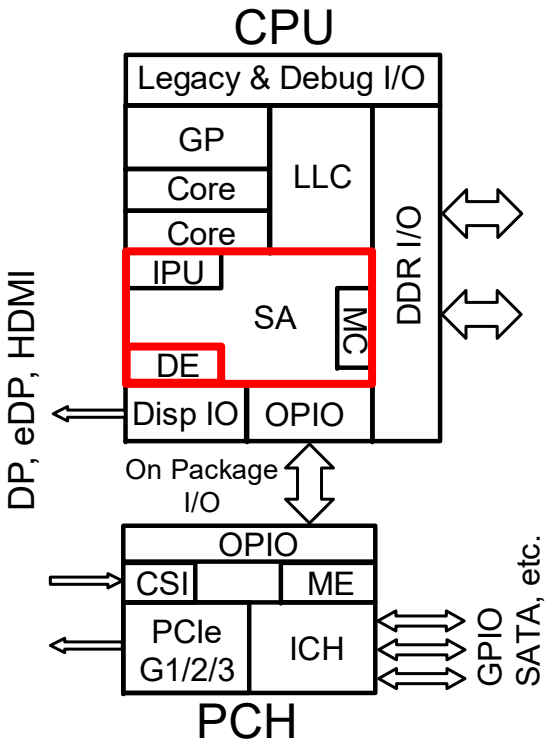

Die floor plan & block diagram (2+2)

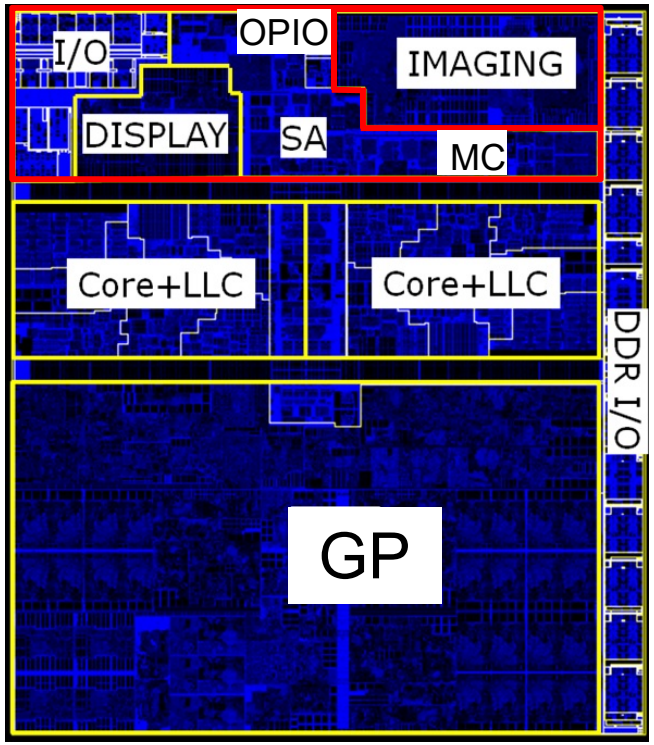

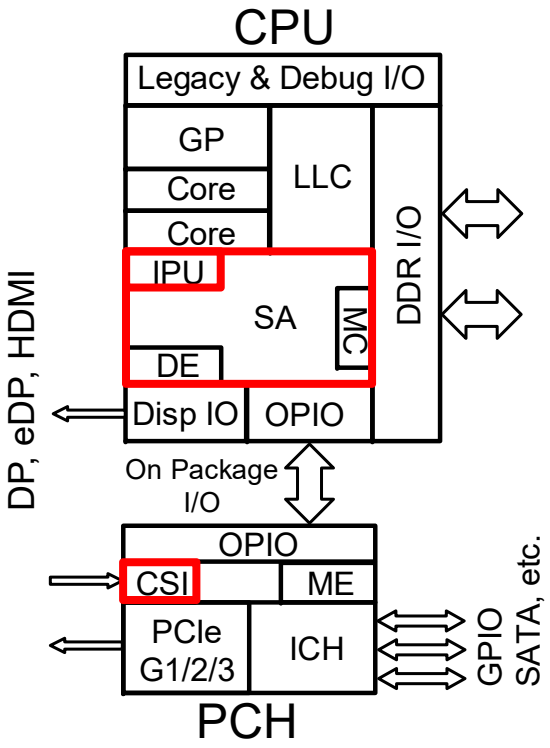

Die floor plan & block diagram (2+2)

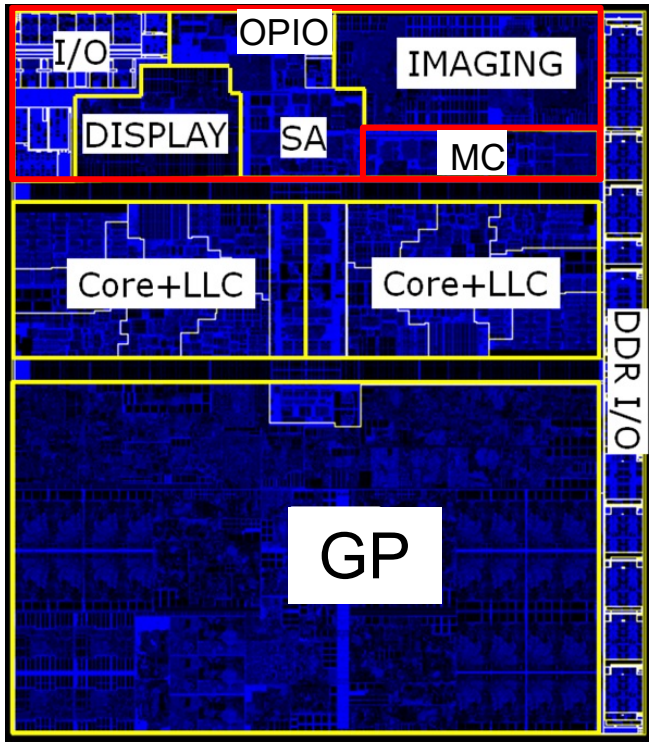

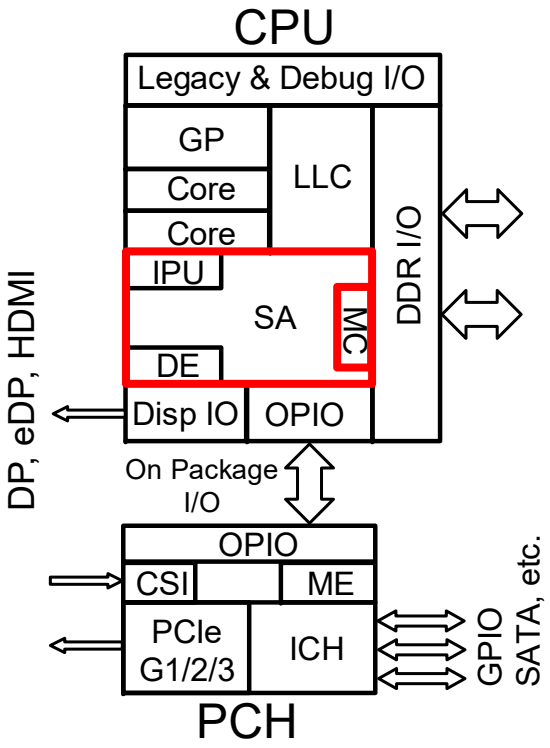

Die floor plan & block diagram (2+2)

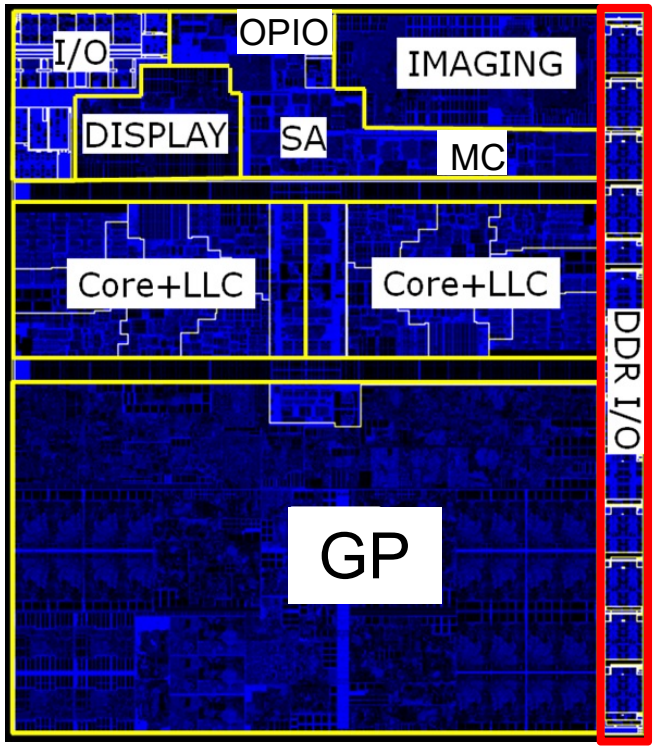

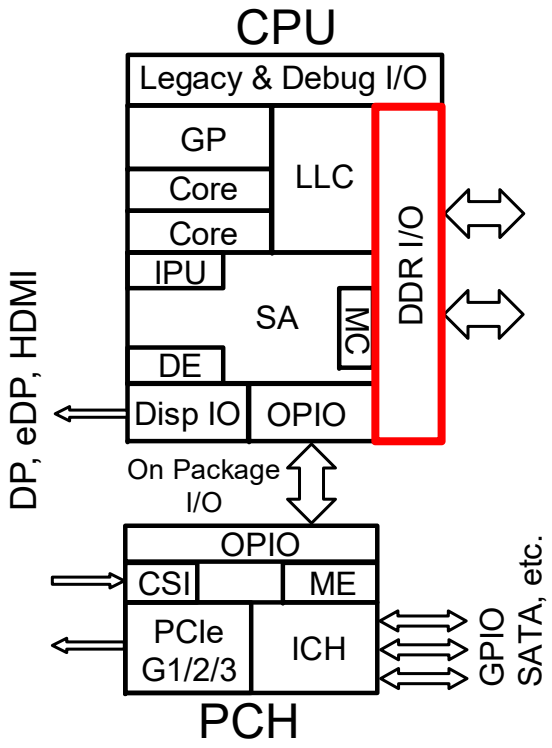

ULT/ULX:
2 channels
64bits/channel
1 – 2.6 GT/s/pin
In DT:
6bits for ECC
Support 4 ranks

Die floor plan & block diagram (2+2)

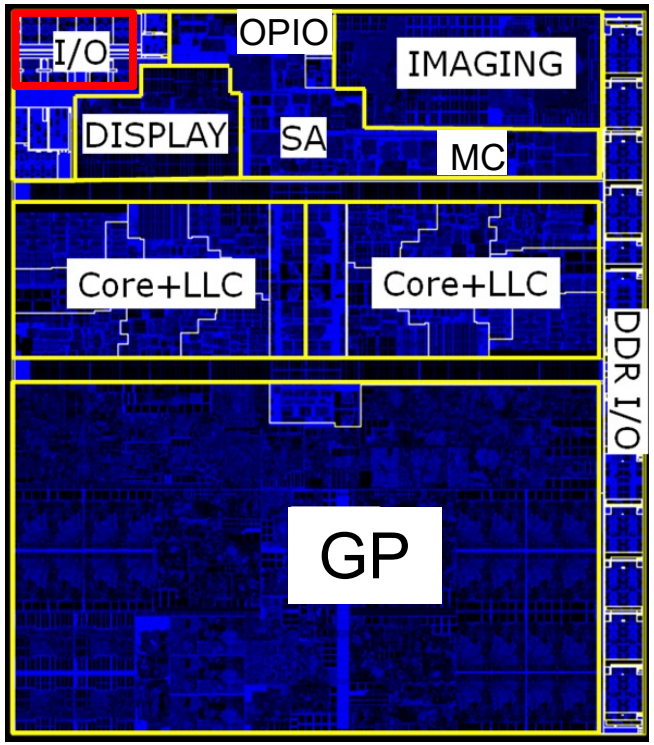

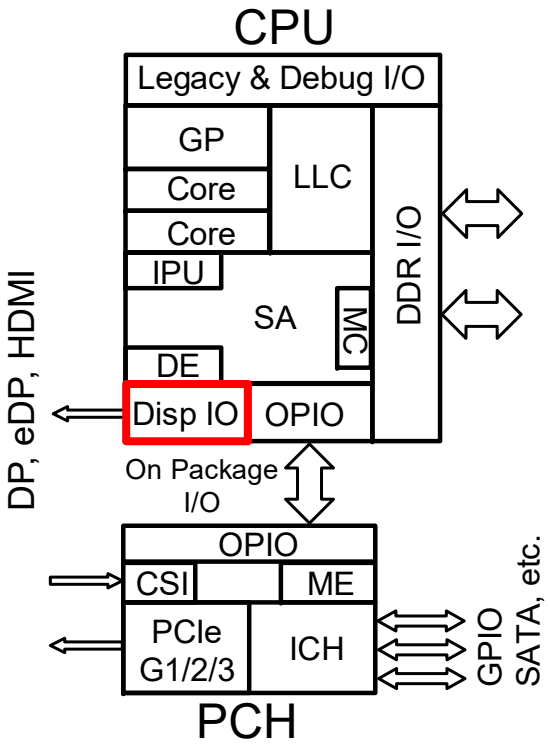

Die floor plan & block diagram (2+2)

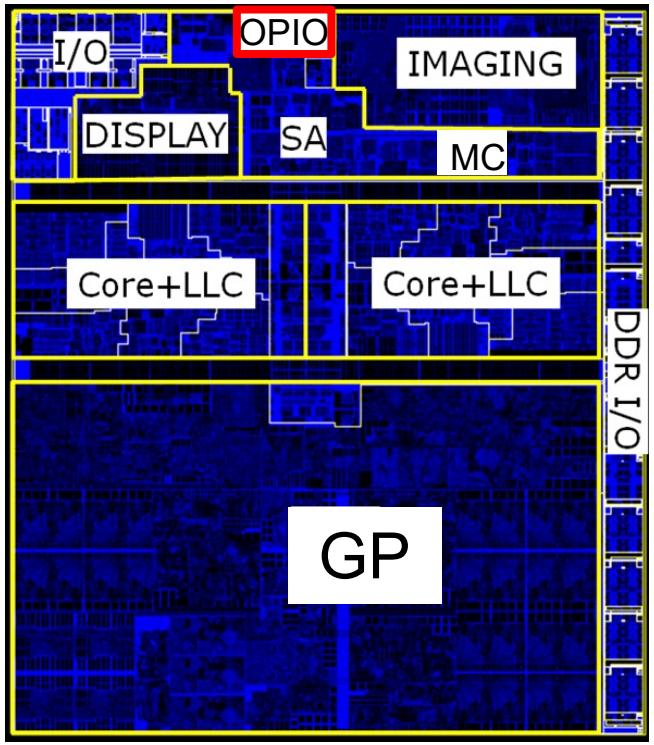

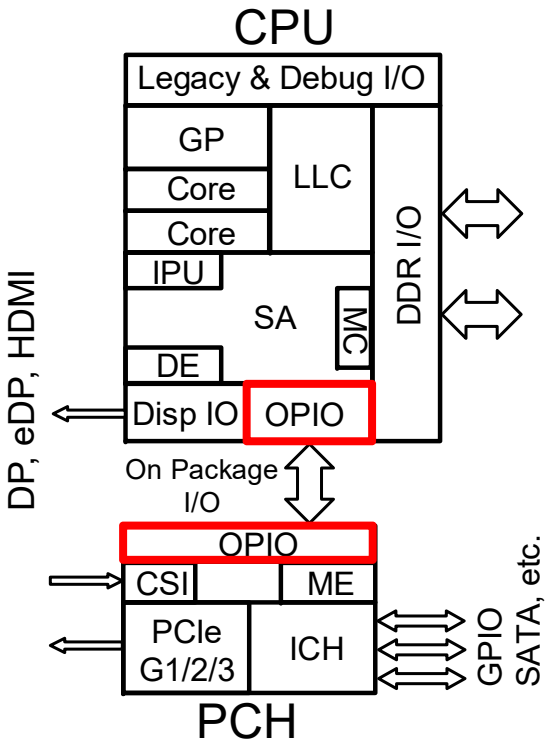

OPIO:
16Tx + 16Rx
4GT/s/pin

Die floor plan & block diagram (2+2)

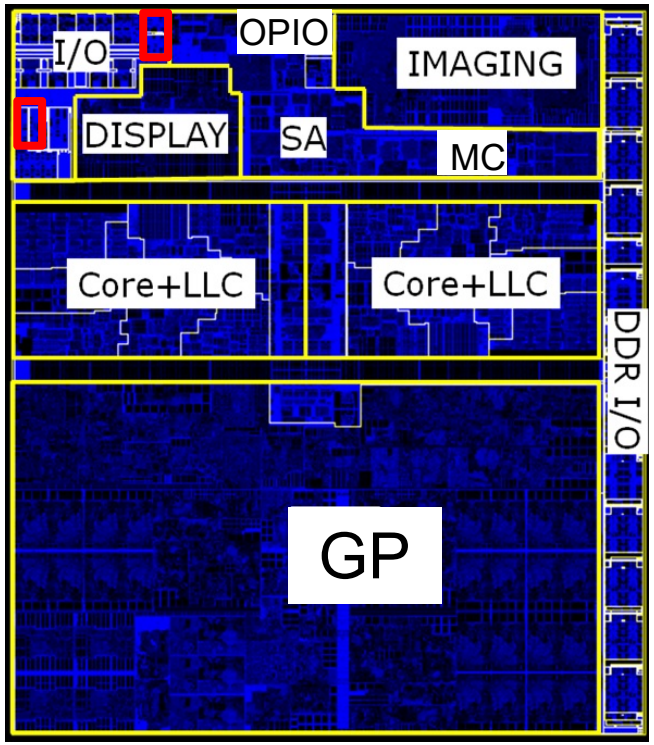

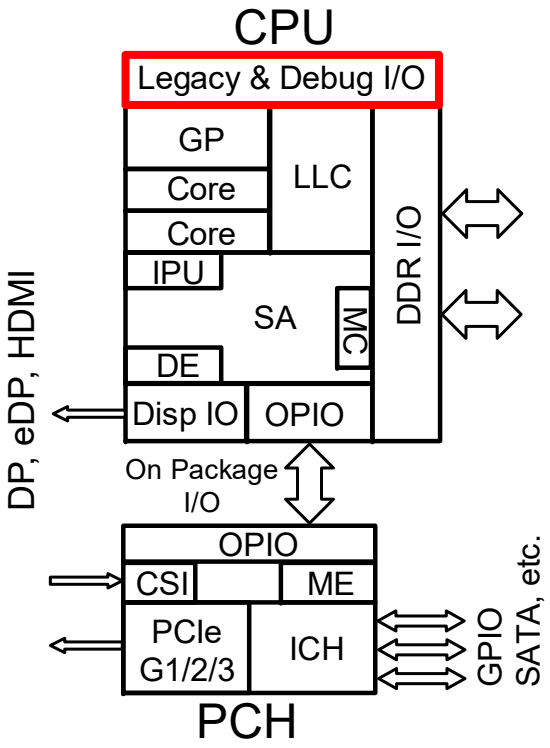

SKL Power Delivery Network (PDN) Topology

- SKL CPU power consumption range is:
  less than 10mW @ standby to higher than 15W @ turbo
- To enable CPU low power states the PDN supports extensive usage of Multi Power Planes (MPP)

SKL Power Delivery Network (PDN) Topology

- SKL CPU power consumption range is:
  less than 10mW @ standby to higher than 15W @ turbo
- To enable CPU low power states the PDN supports extensive usage of Multi Power Planes (MPP)
  - Using fine granularity, low latency, Power Gates (PGT)

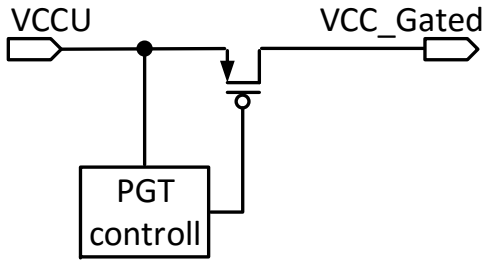

SKL Power Delivery Network (PDN) Topology

- SKL CPU power consumption range is:
  less than 10mW @ standby to higher than 15W @ turbo
- To enable CPU low power states the PDN supports extensive usage of Multi Power Planes (MPP)
  - Using fine granularity, low latency, Power Gates (PGT)
  - Using local data retention, based on Power MUXing

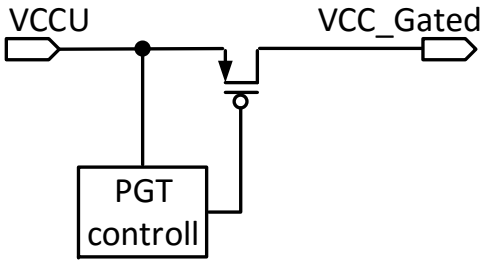

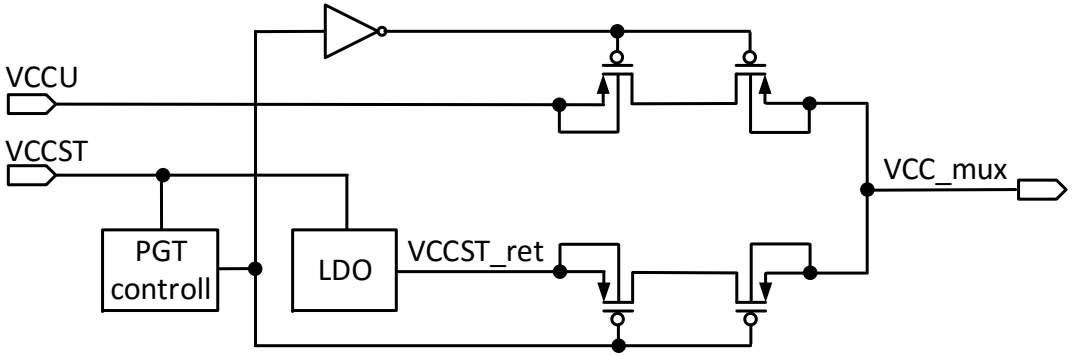

IA CORE PDN breakdown for power efficiency

- PGT disable non-functional logic for low Thermal Design Power (TDP)
  - Main PGT shut off the entire IA CORE

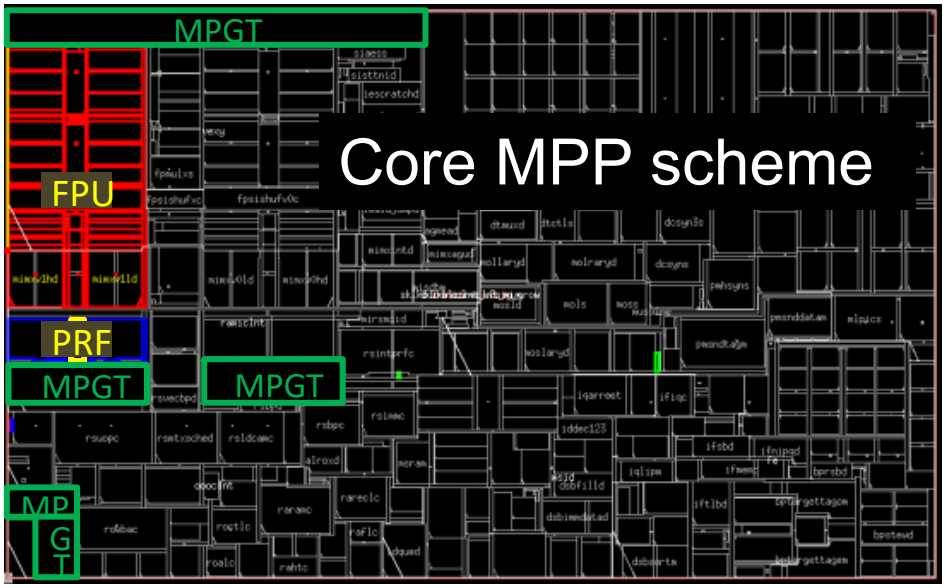

IA CORE PDN breakdown for power efficiency

- PGT disable non-functional logic for low Thermal Design Power (TDP)
  - Main PGT shut off the entire IA CORE
  - EXE unit FPU & PRF can switched OFF separately by a 2nd level PGT

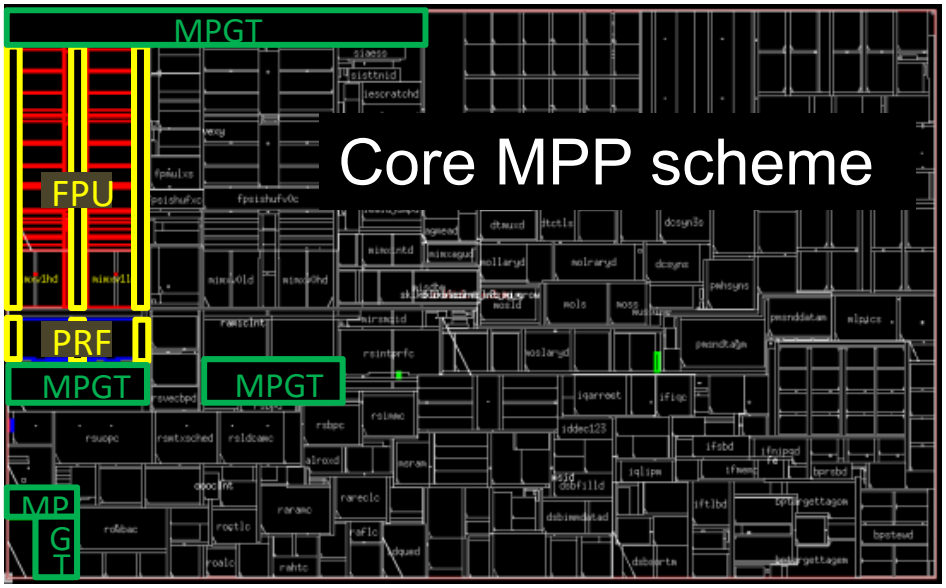

SA PDN breakdown for power efficiency

- Disable non-functional logic for low energy states (SA remains “ON”)
  - 4 independent power domains represents by the different colors

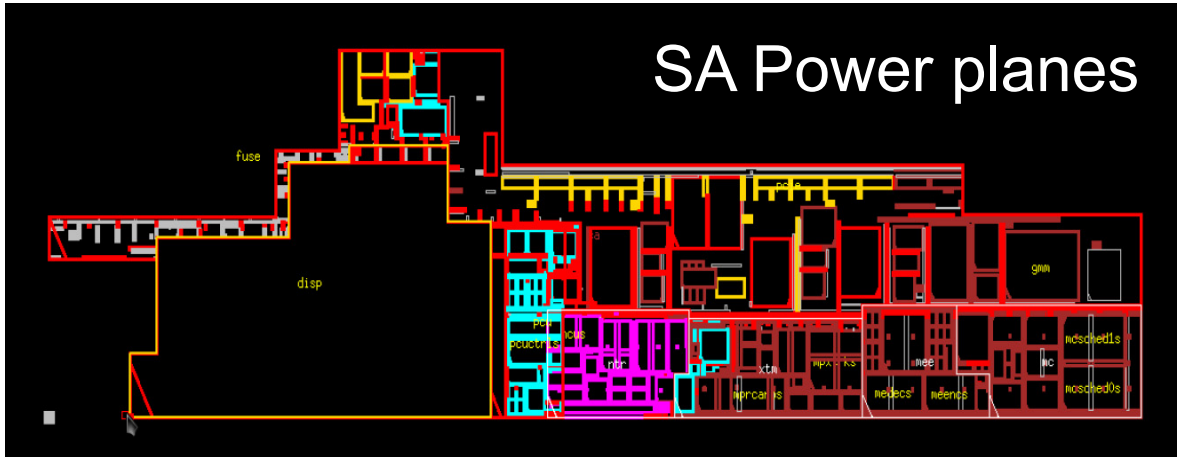

SA PDN breakdown for power efficiency

- Disable non-functional logic for low energy states (SA remains “ON”)
  - 4 independent power domains represents by the different colors
  - 3 retention domains for ultra-low power (including Power MUXing)

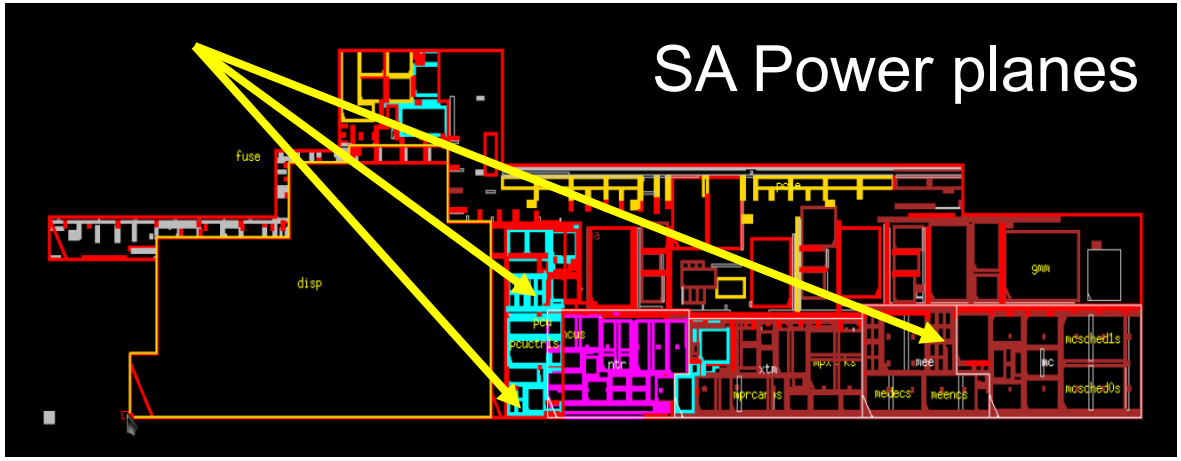

PGT power-up

- Enable fast power up while meeting reliability goals

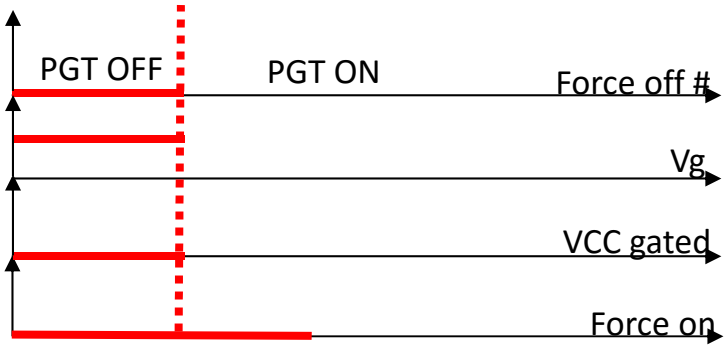

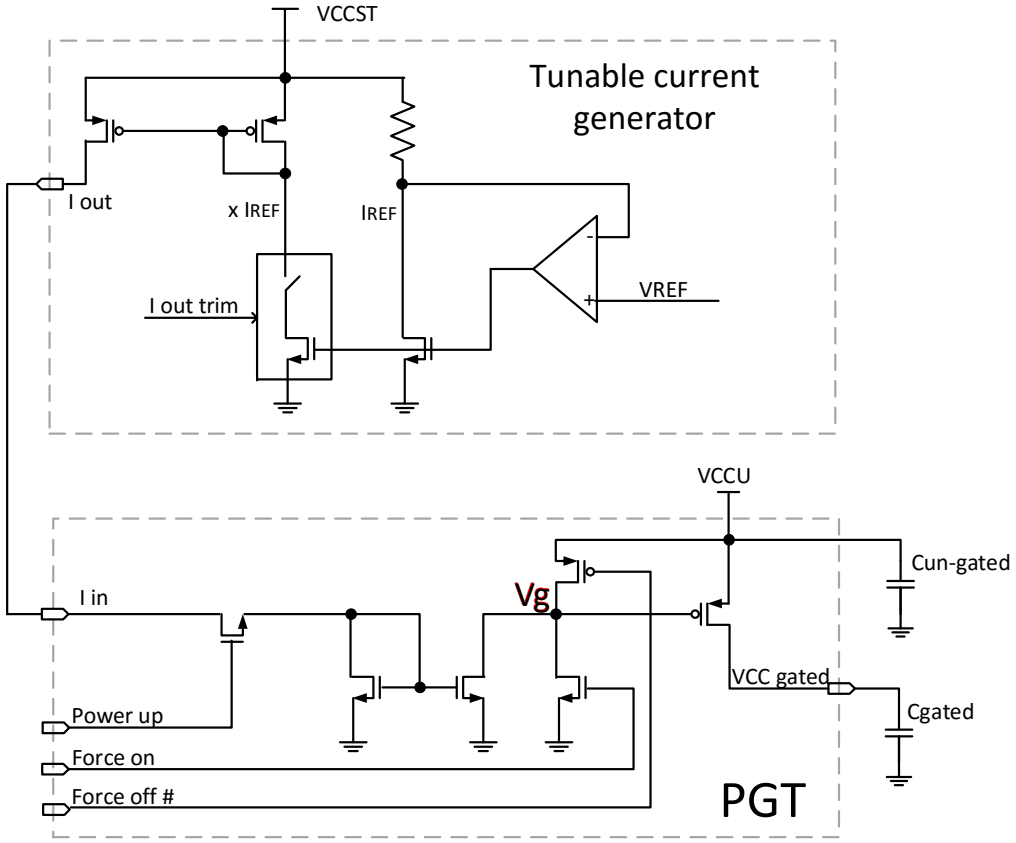

PGT power-up

- Enable fast power up while meeting reliability goals

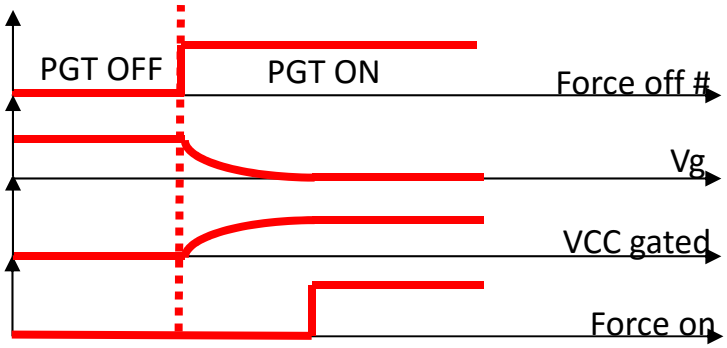

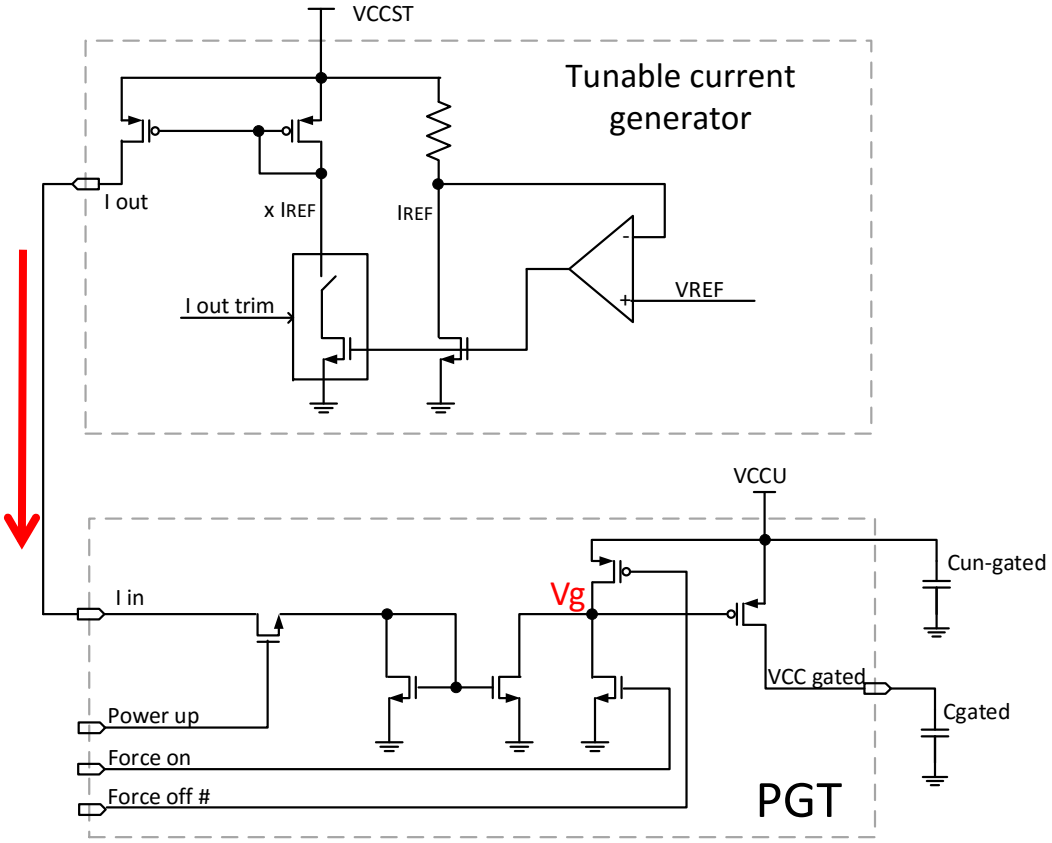

PGT power-up

- Enable fast power up while meeting reliability goals

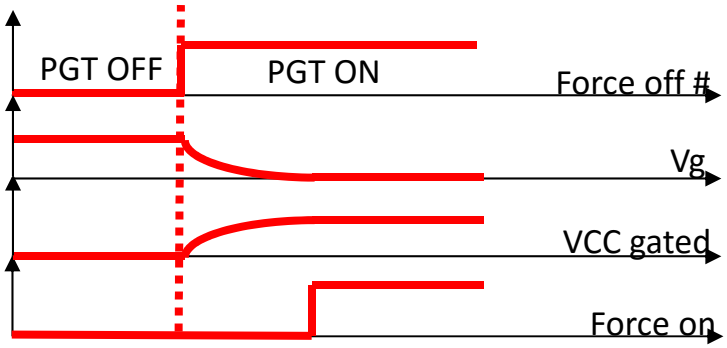

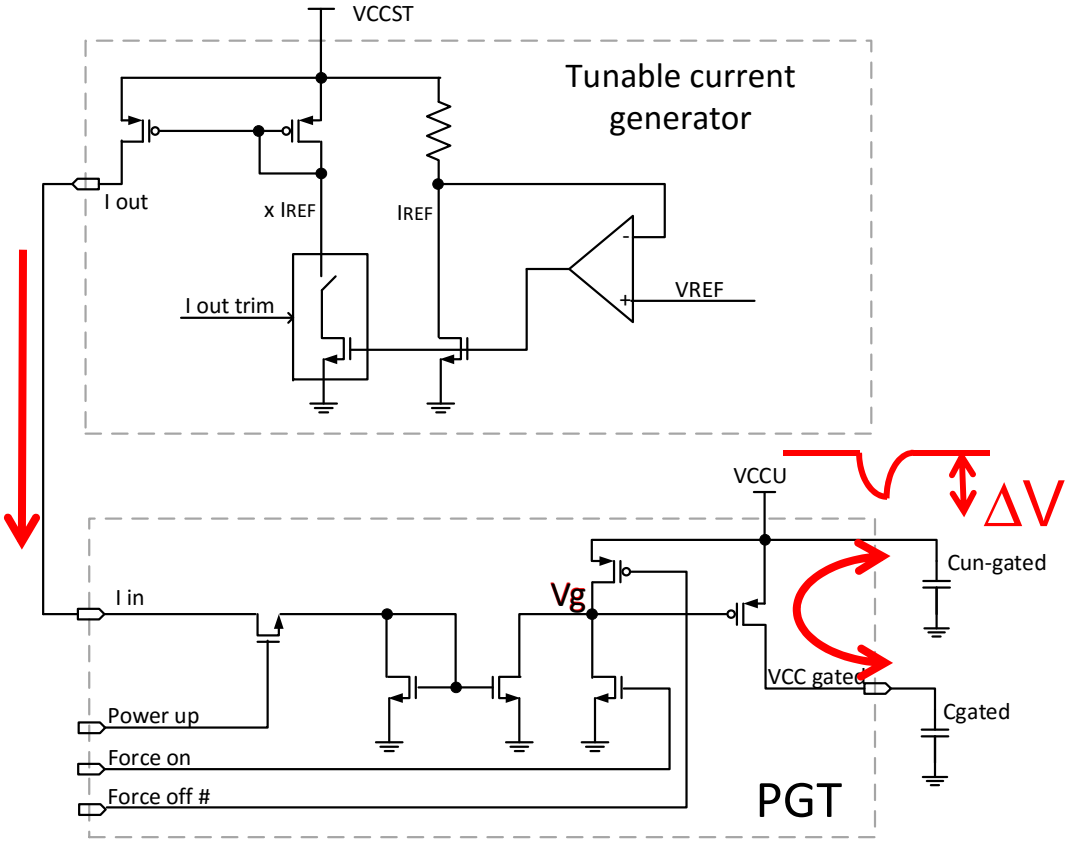

PGT power-up

- Enable fast power up while meeting reliability goals
- Allows post Si trimming to optimize droop on VCCU vs. power-up speed (droop is proportional to Cgate/Cun-gated)

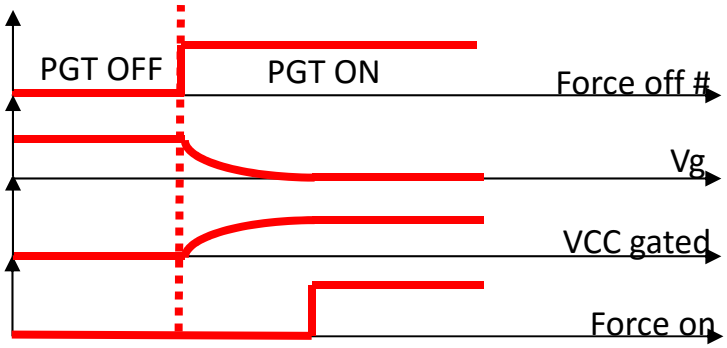

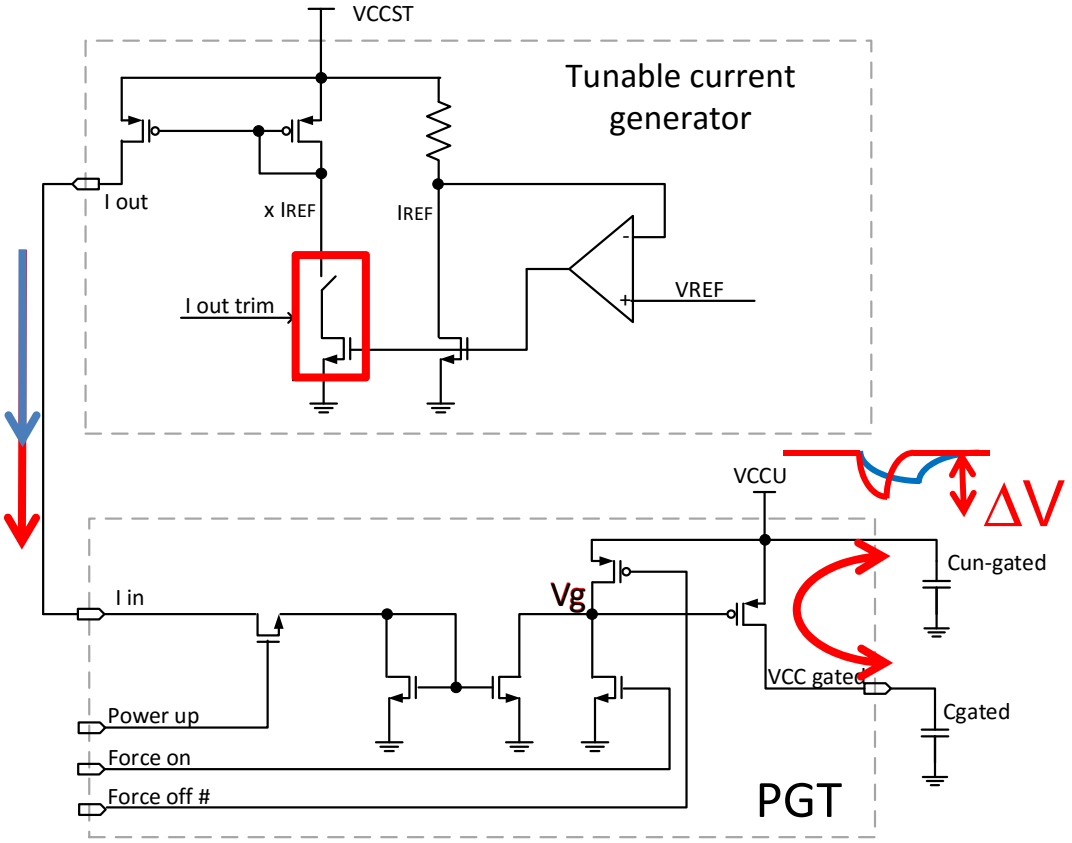

Power Gates Measured Si Performance

EXE
voltage rate @ power up

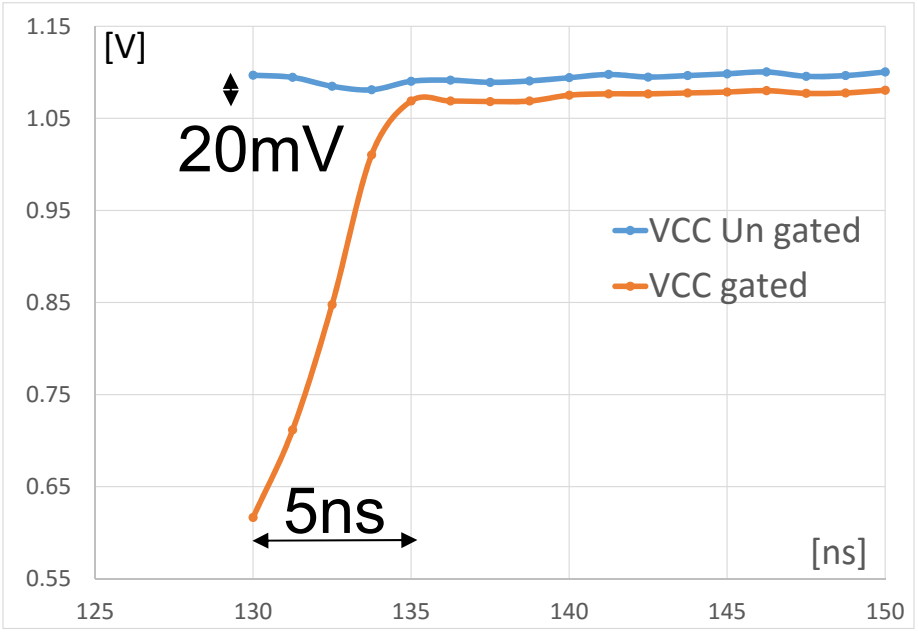

OPIO
voltage droop @ power up

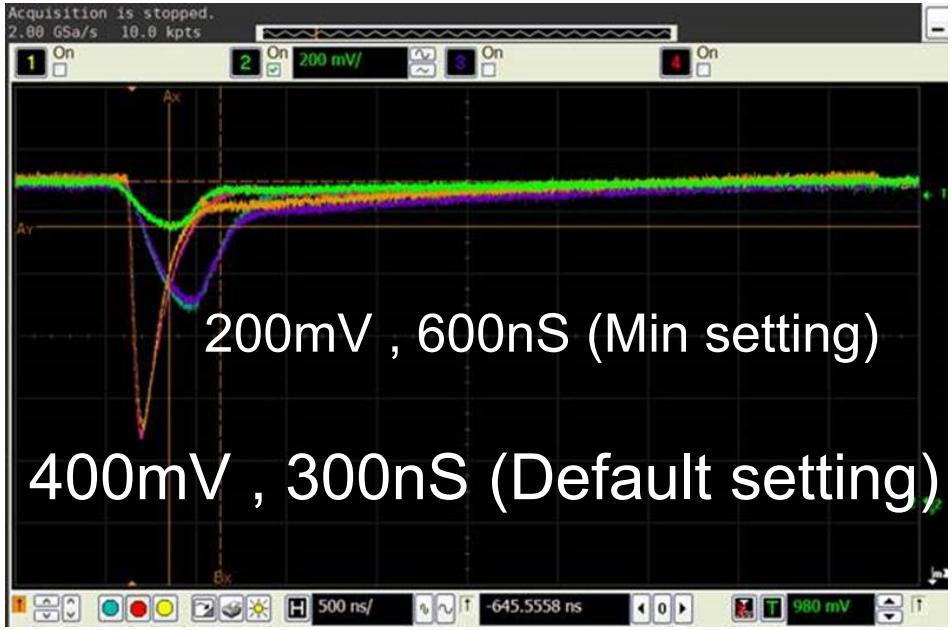

Voltage rate @ PGT power up @ low Cgated (30nf)

Voltage droop vs. PGT power up speed @ high Cgated (4.7uF)

SKL low power clocking

- Digital PLL's implementation in all logic clock domains
- Fractional-n (Fn) PLL implementation in Display IO
- Implementation of Island based Core Clock distribution
- Clock distribution consolidation in the SA
  - Single PLL driving multiple clock spines
  - Single spine driving multiple clock & voltage domains

Digital PLL vs. Analog PLL (RBPLL)

- RBPLL is based on a self-biased replica buffer

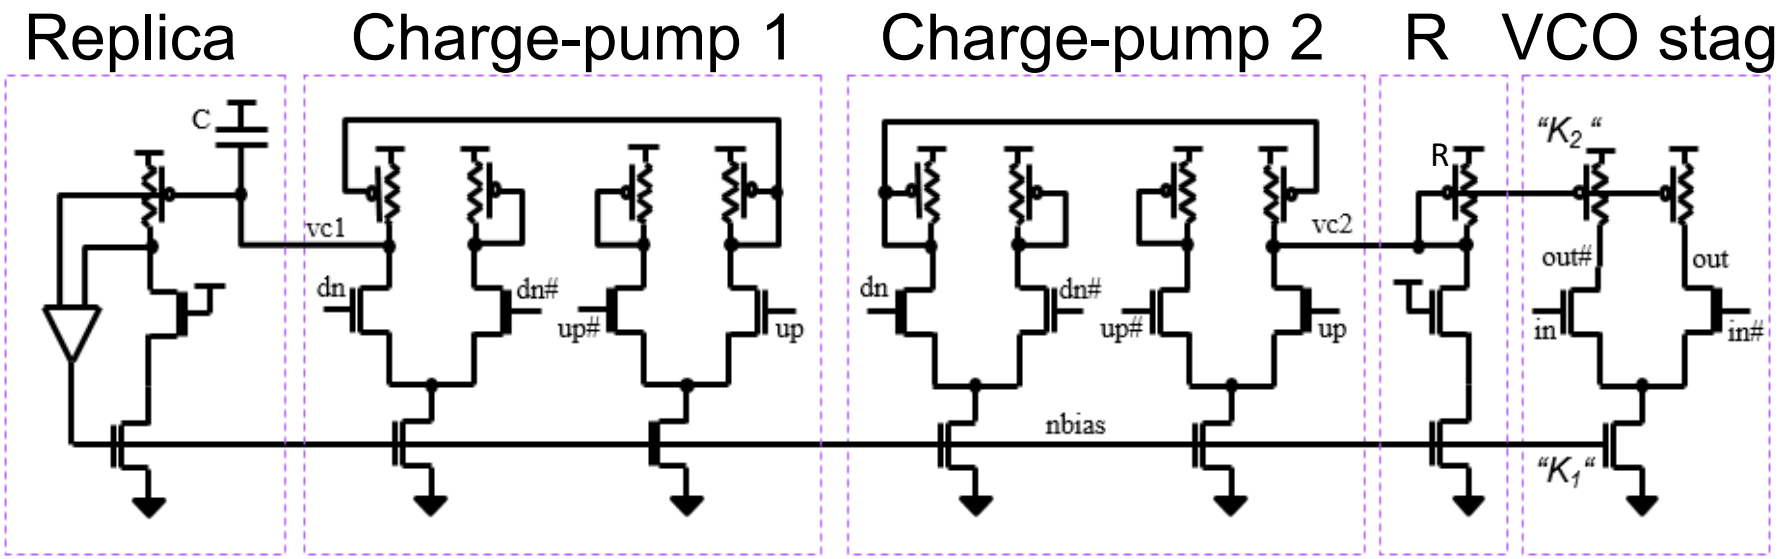

Digital PLL vs. Analog PLL (RBPLL)

- RBPLL is based on a self-biased replica buffer
  - All circuit blocks consume DC current that is switched between the differential stage branches

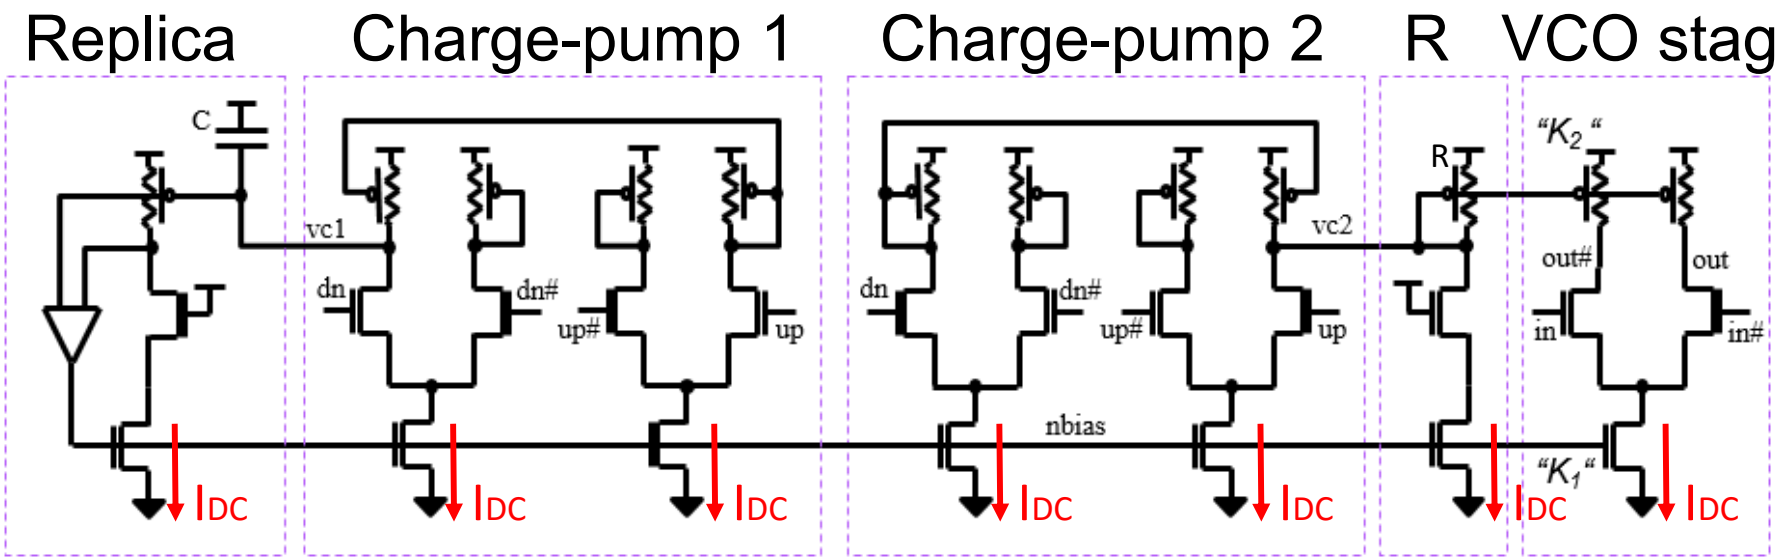

Digital PLL ROSC DCO stage

- CMOS stage replace a DC biased CML stage

Digital PLL DCO stage

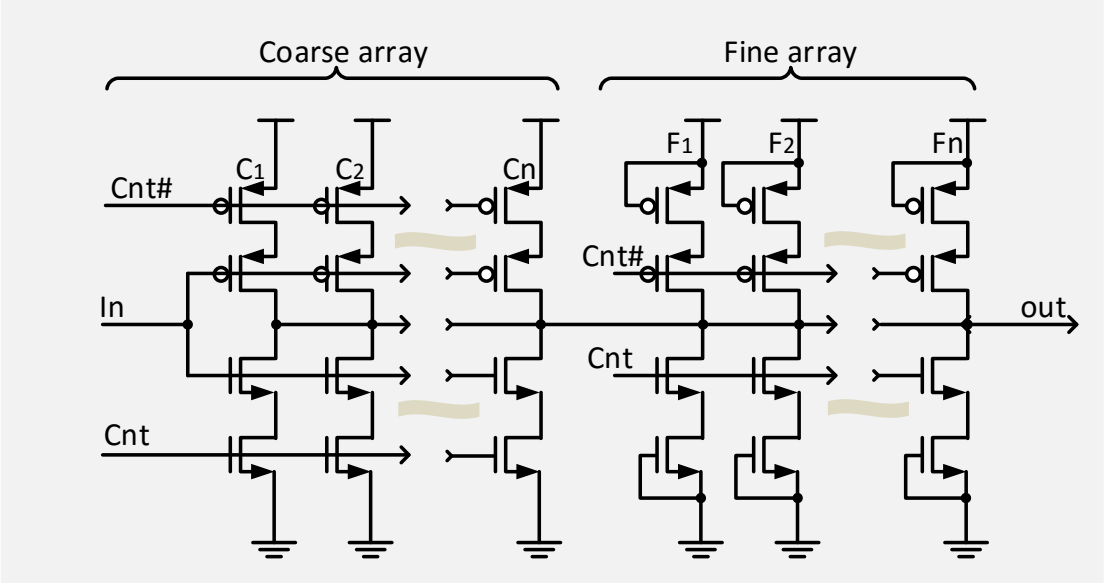

VCO stage + bias

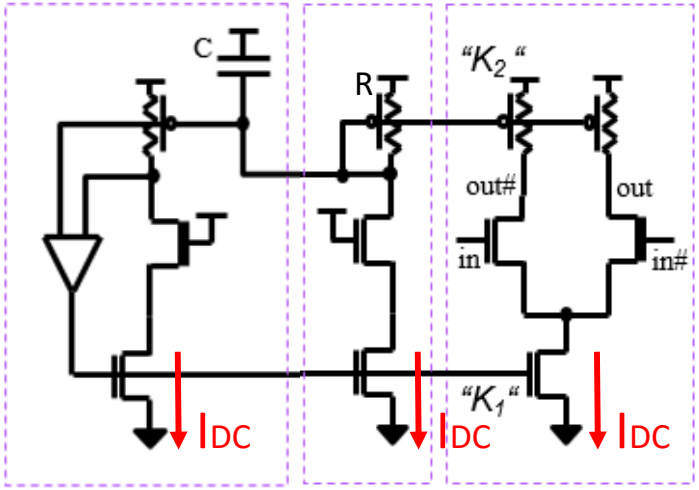

Digital PLL ROSC DCO stage

- CMOS stage replace a DC biased CML stage

Digital PLL DCO stage

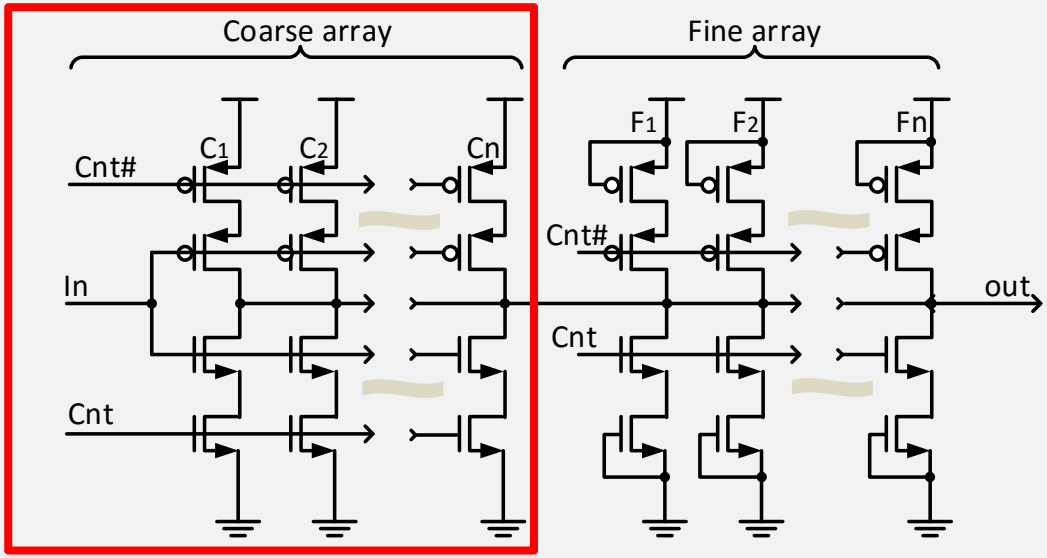

VCO stage + bias

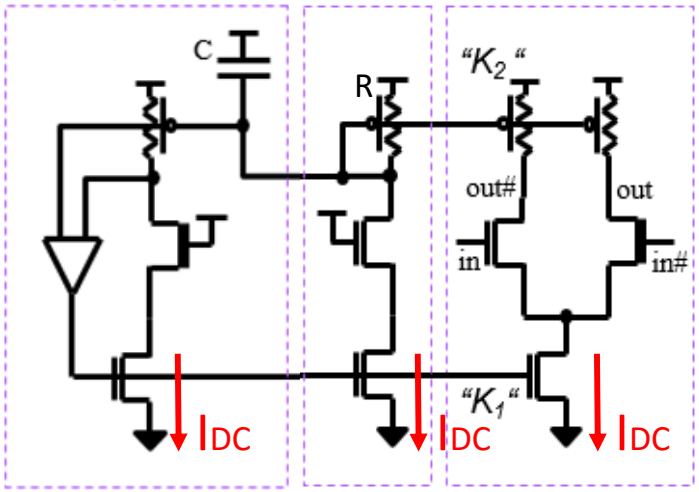

Digital PLL ROSC DCO stage

- CMOS stage replace a DC biased CML stage

Digital PLL DCO stage

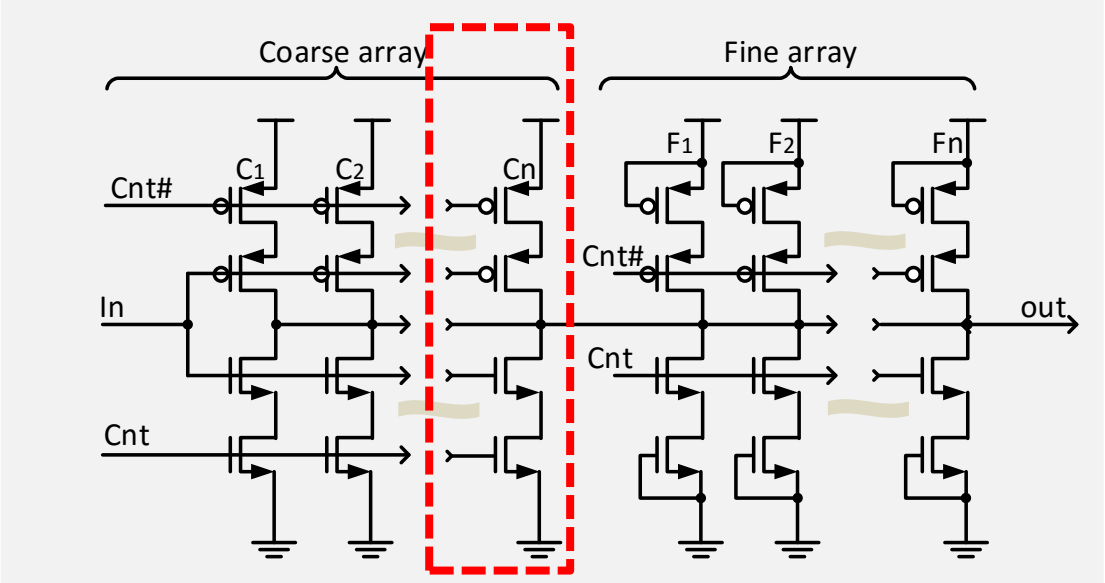

VCO stage + bias

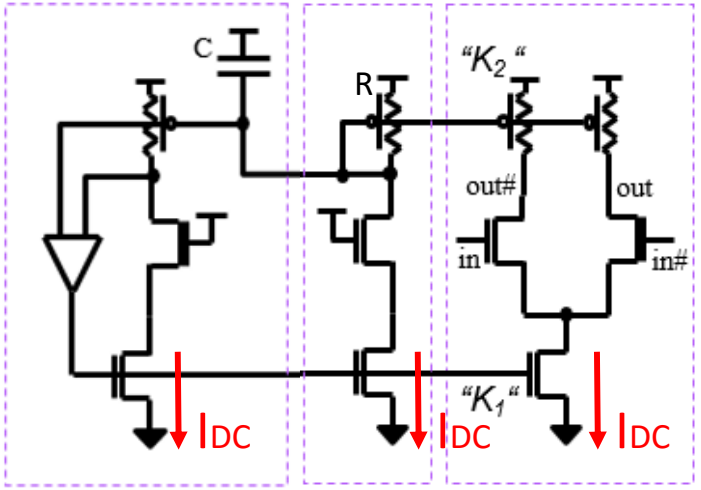

Digital PLL ROSC DCO stage

- CMOS stage replace a DC biased CML stage

Digital PLL DCO stage

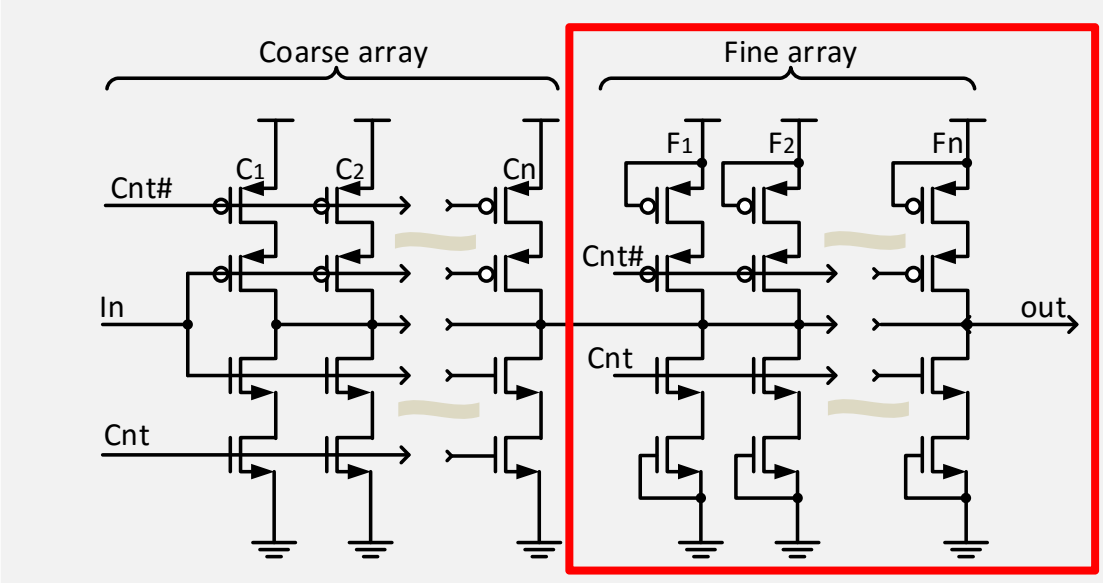

VCO stage + bias

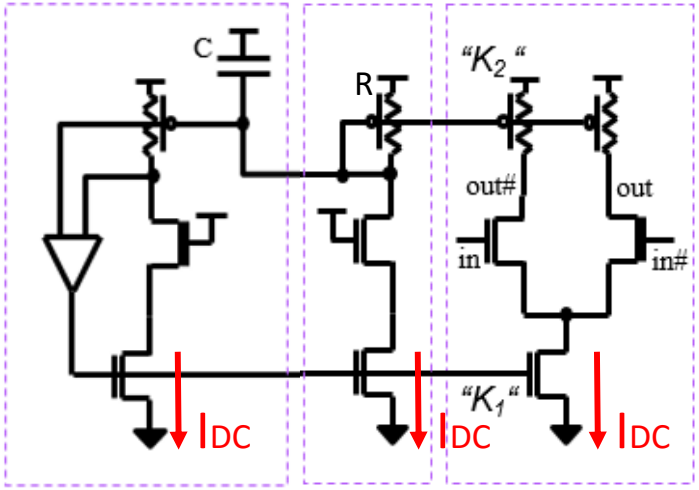

Digital PLL ROSC DCO stage

- CMOS stage replace a DC biased CML stage

Digital PLL DCO stage

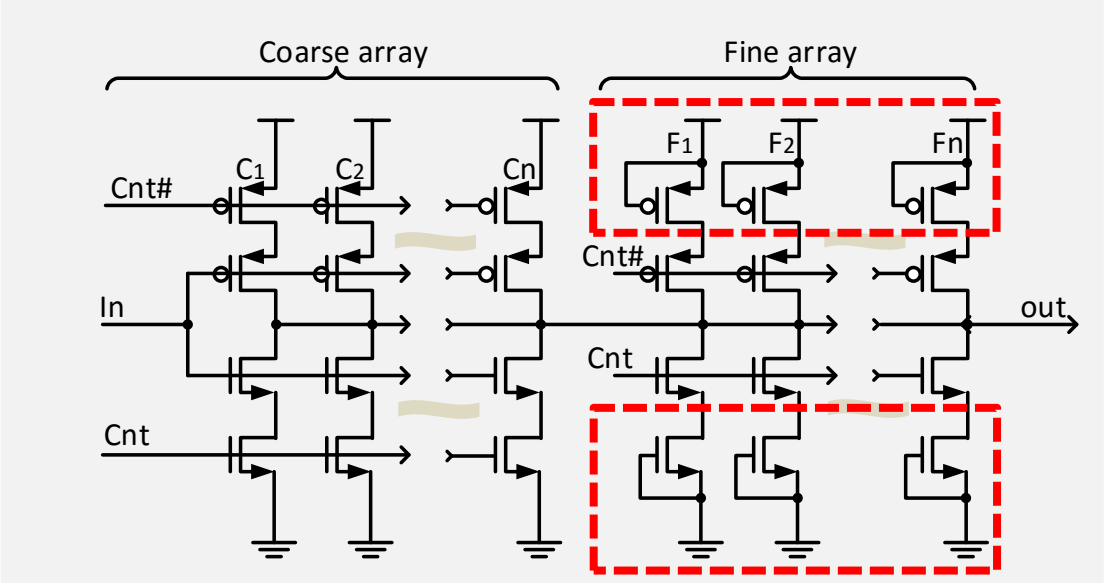

VCO stage + bias

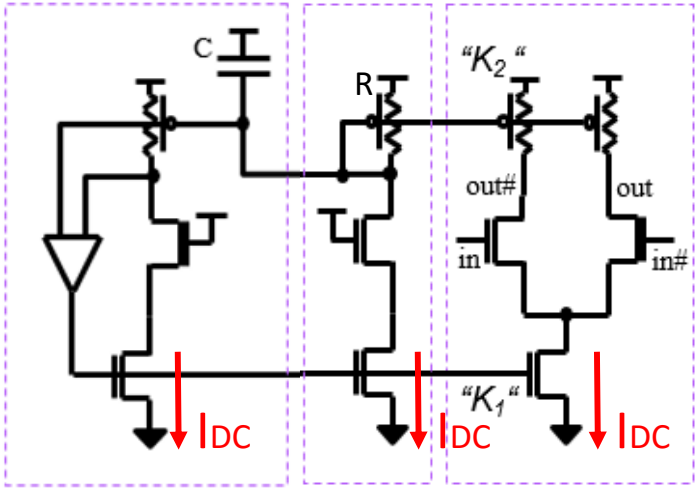

Digital PLL Phase detector (TDC) & DLF

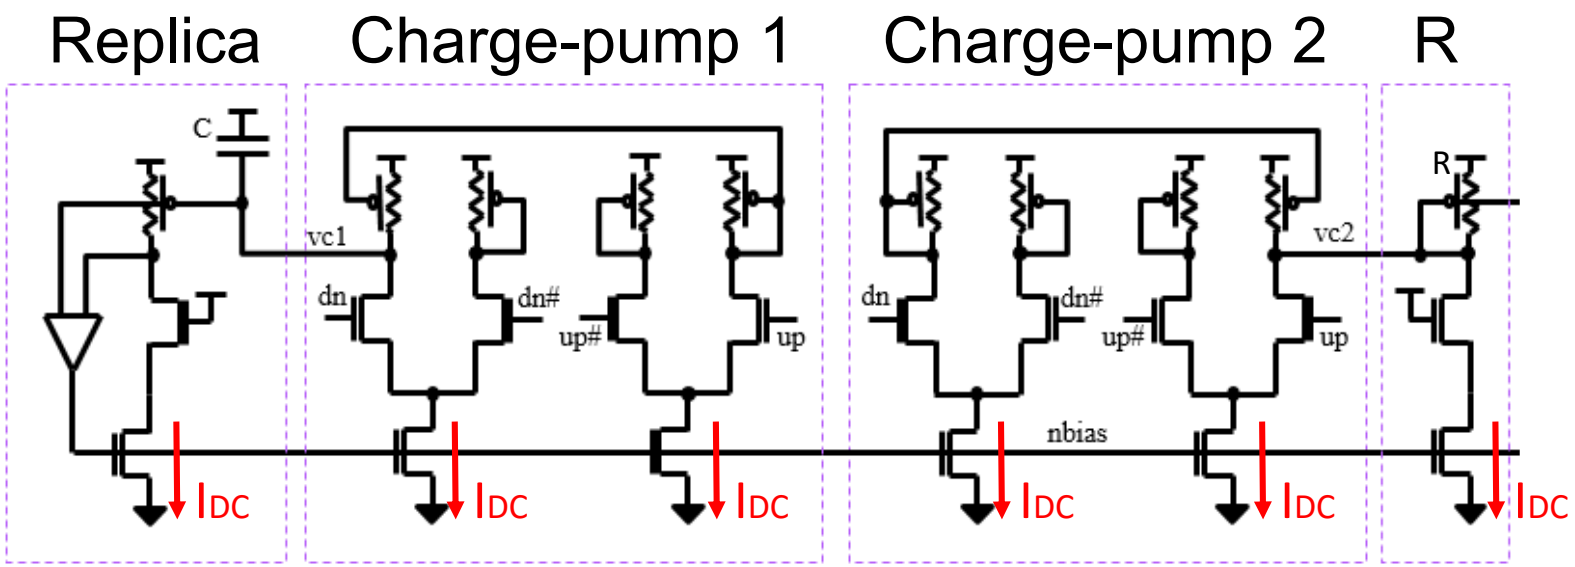

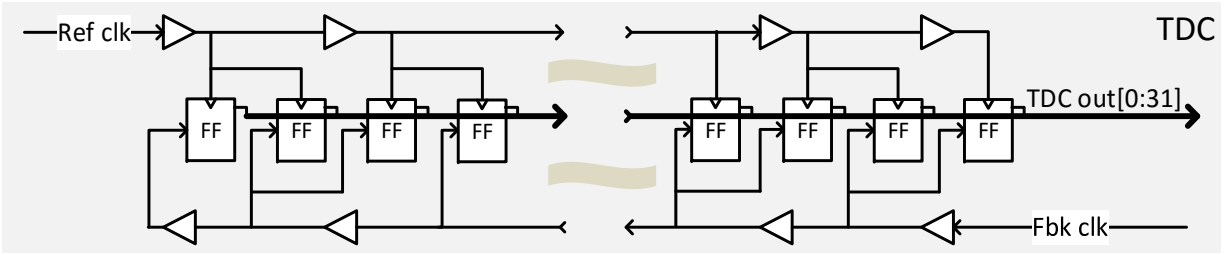

TDC

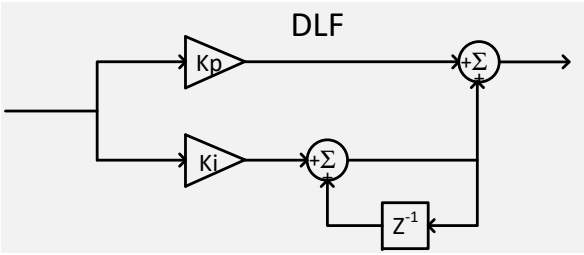

DLF

Digital PLL Phase detector (TDC) & DLF

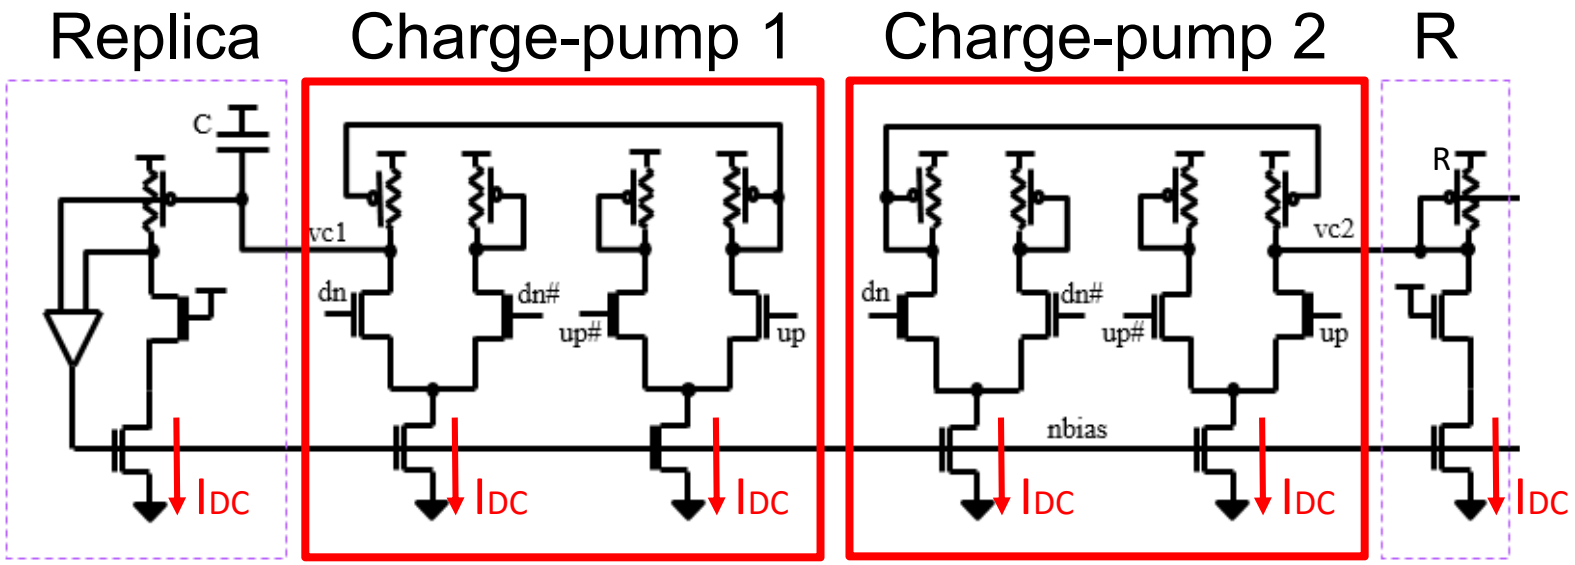

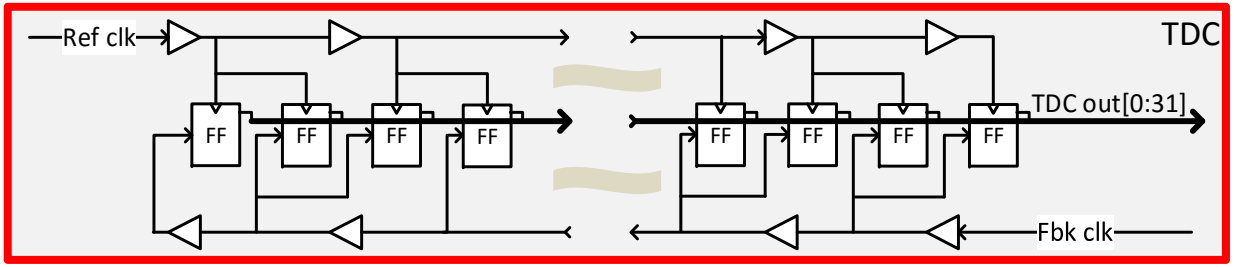

TDC

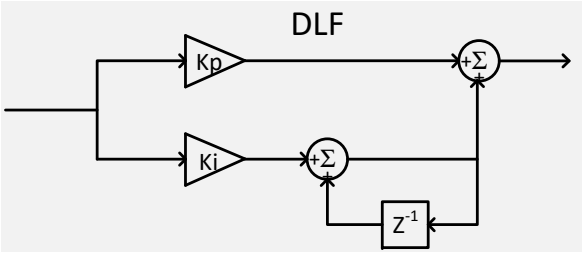

DLF

Digital PLL Phase detector (TDC) & DLF

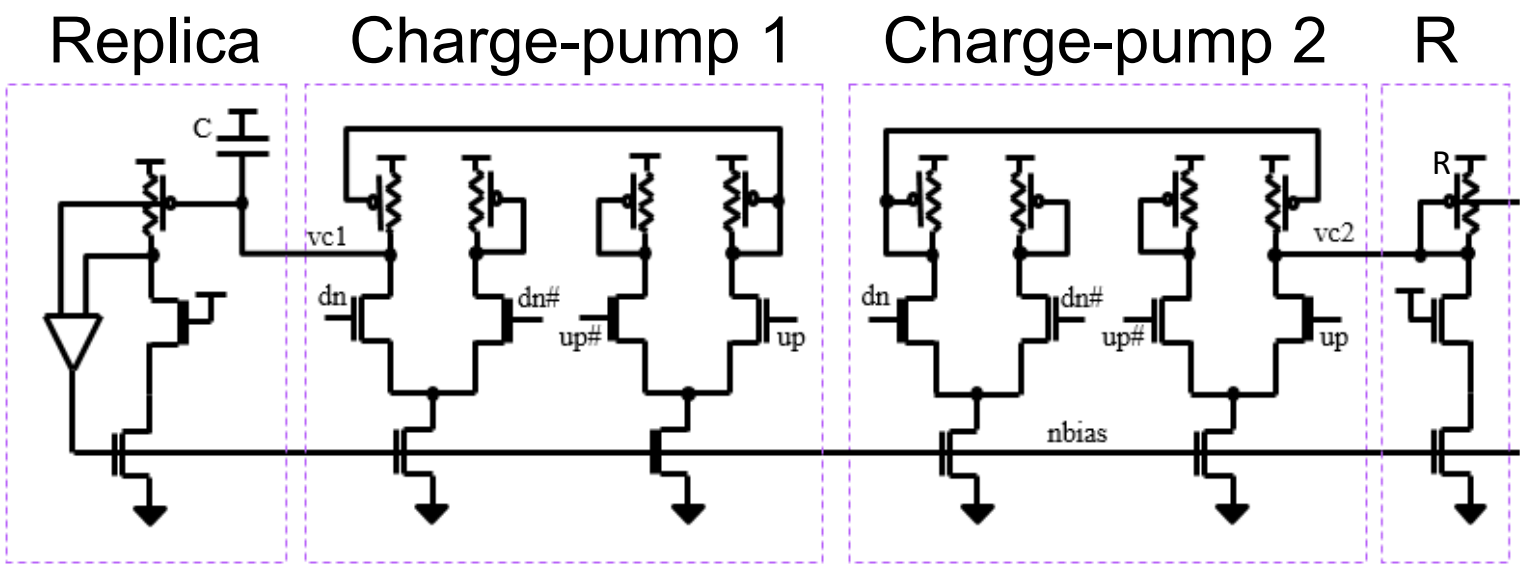

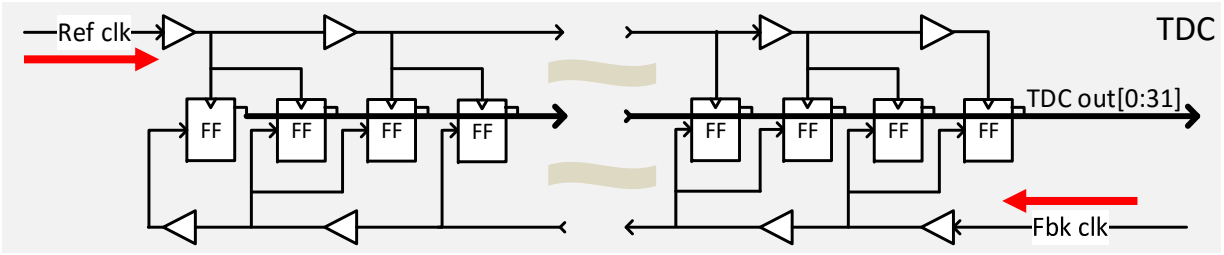

TDC

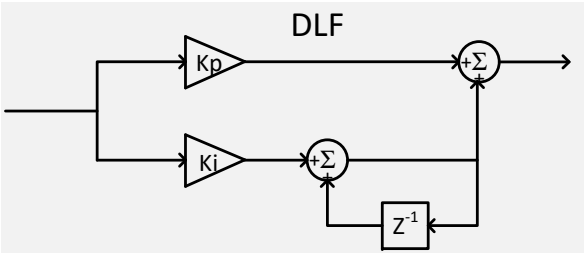

DLF

Digital PLL Phase detector (TDC) & DLF

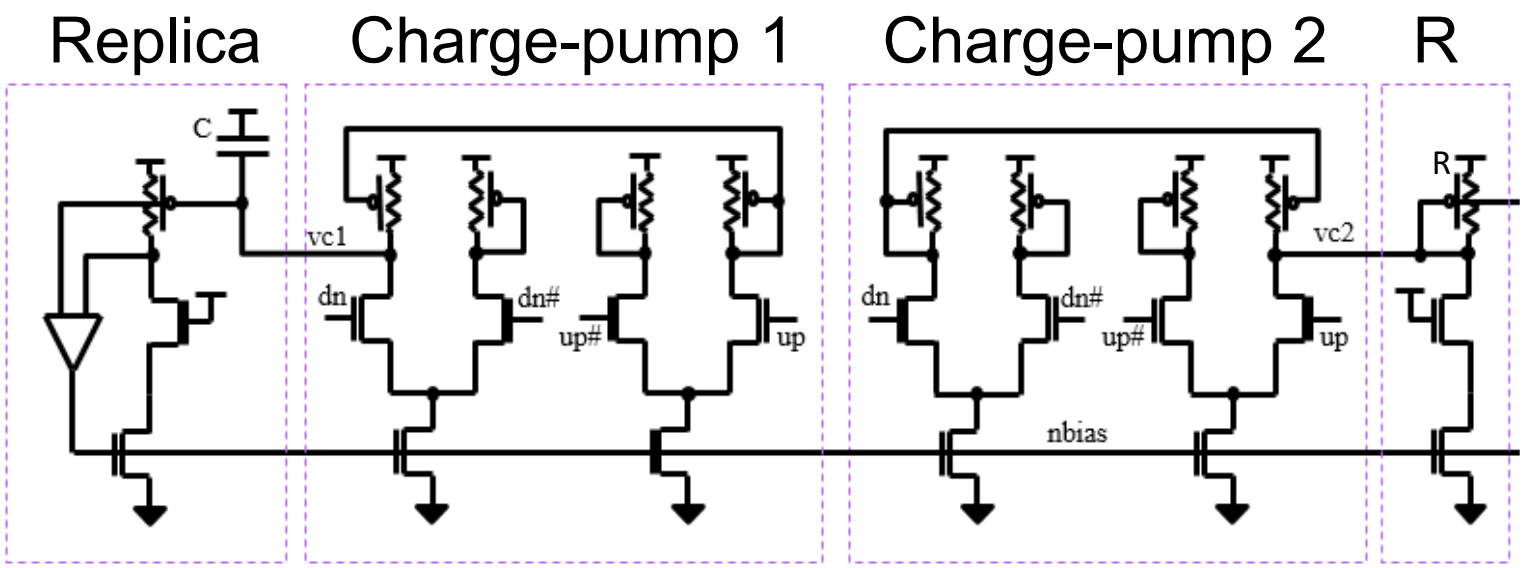

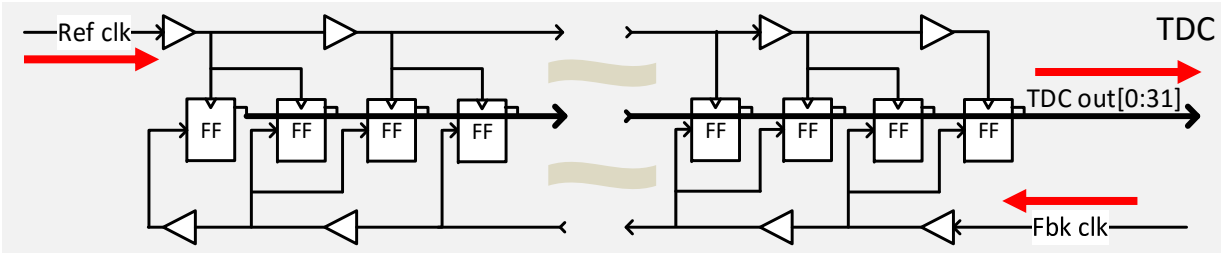

TDC

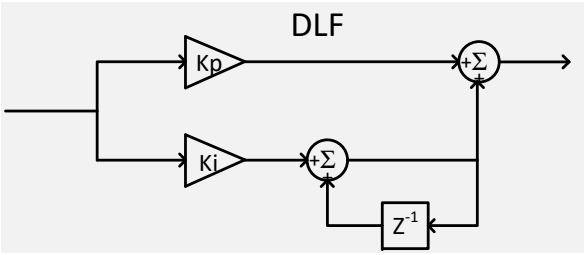

DLF

Digital PLL Phase detector (TDC) & DLF

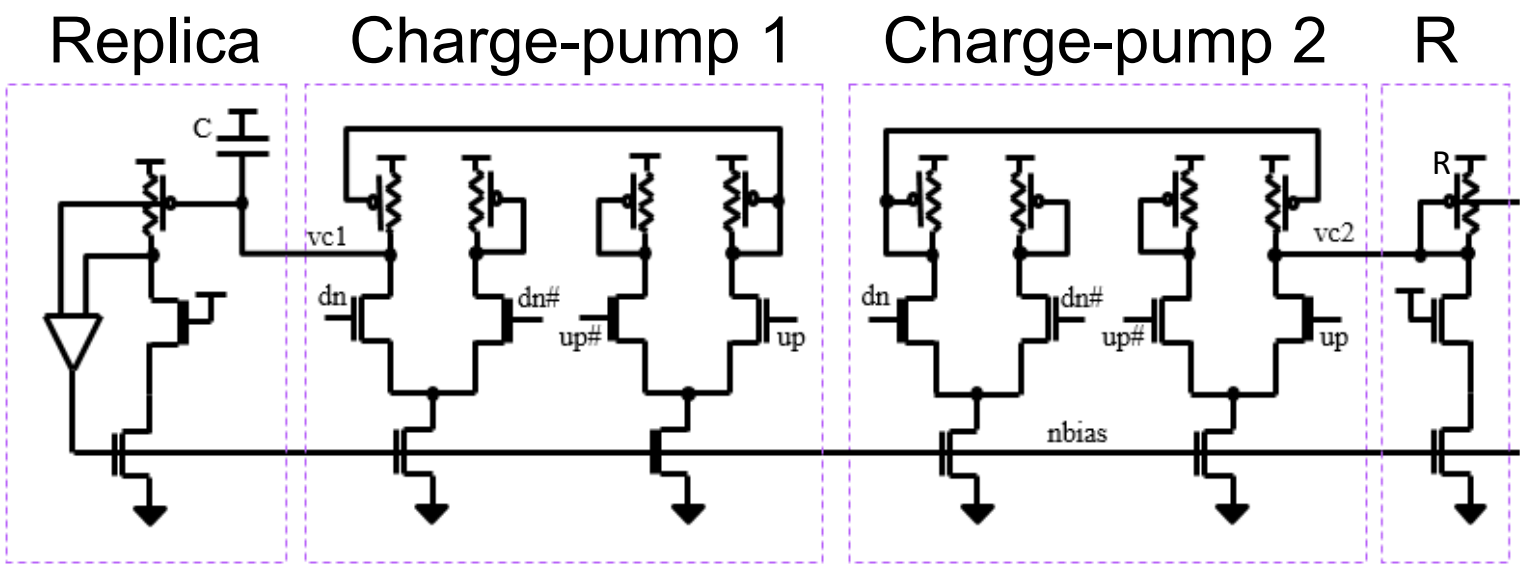

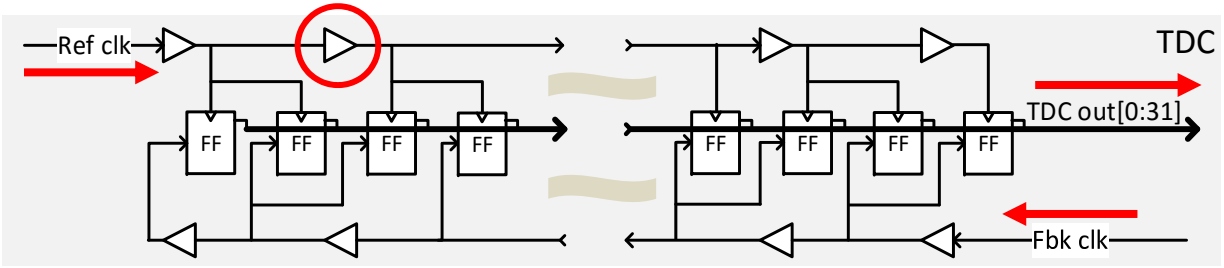

TDC

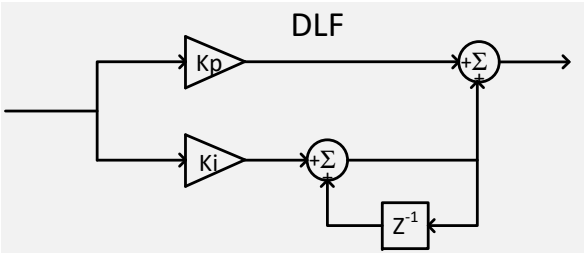

DLF

Digital PLL Phase detector (TDC) & DLF

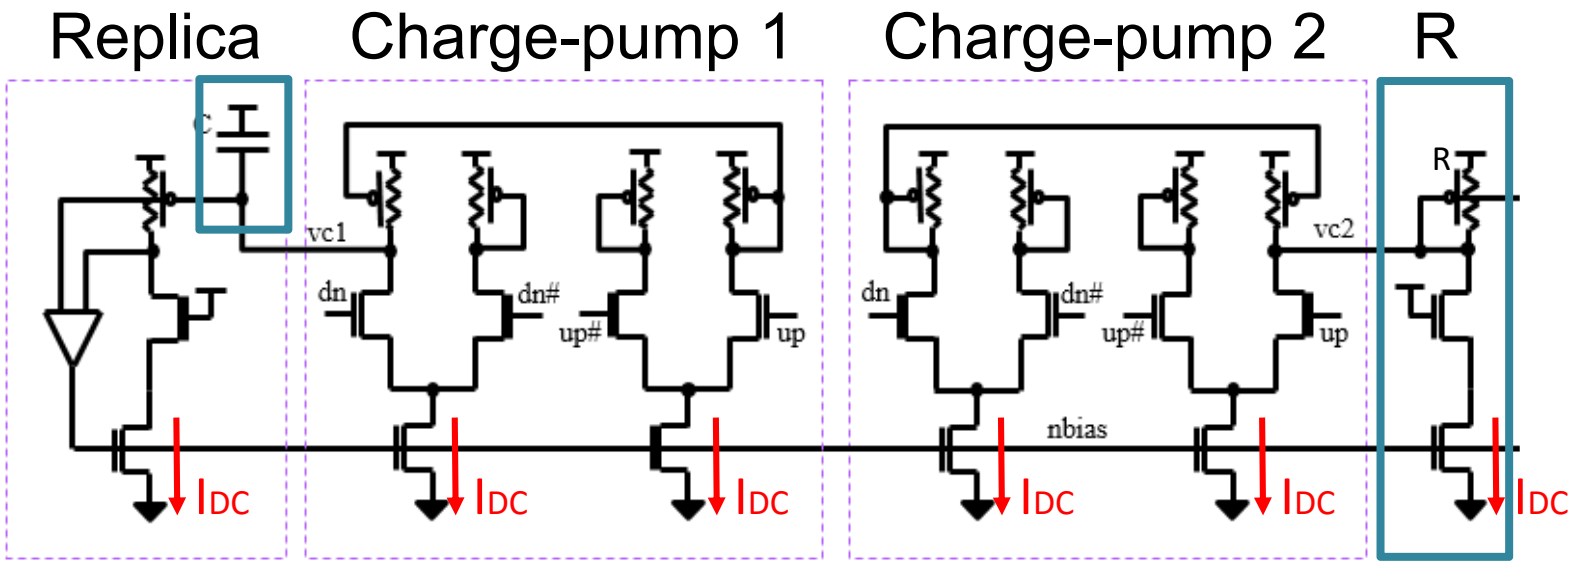

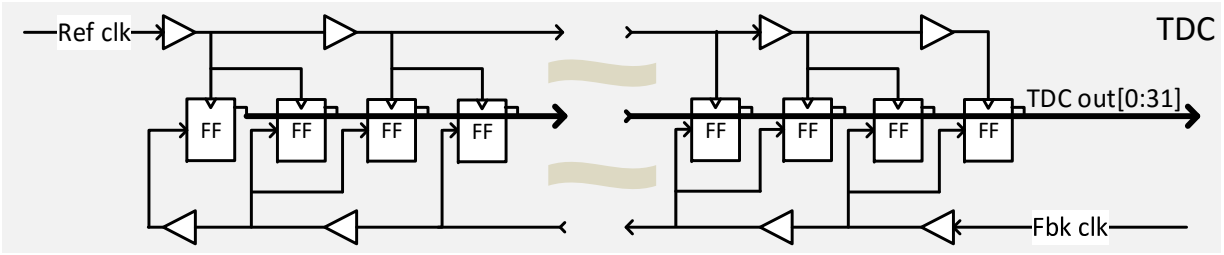

TDC

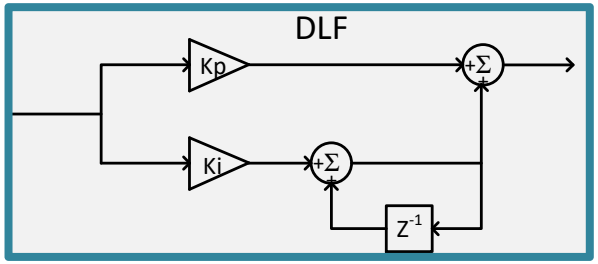

DLF

Digital PLL Phase detector (TDC) & DLF

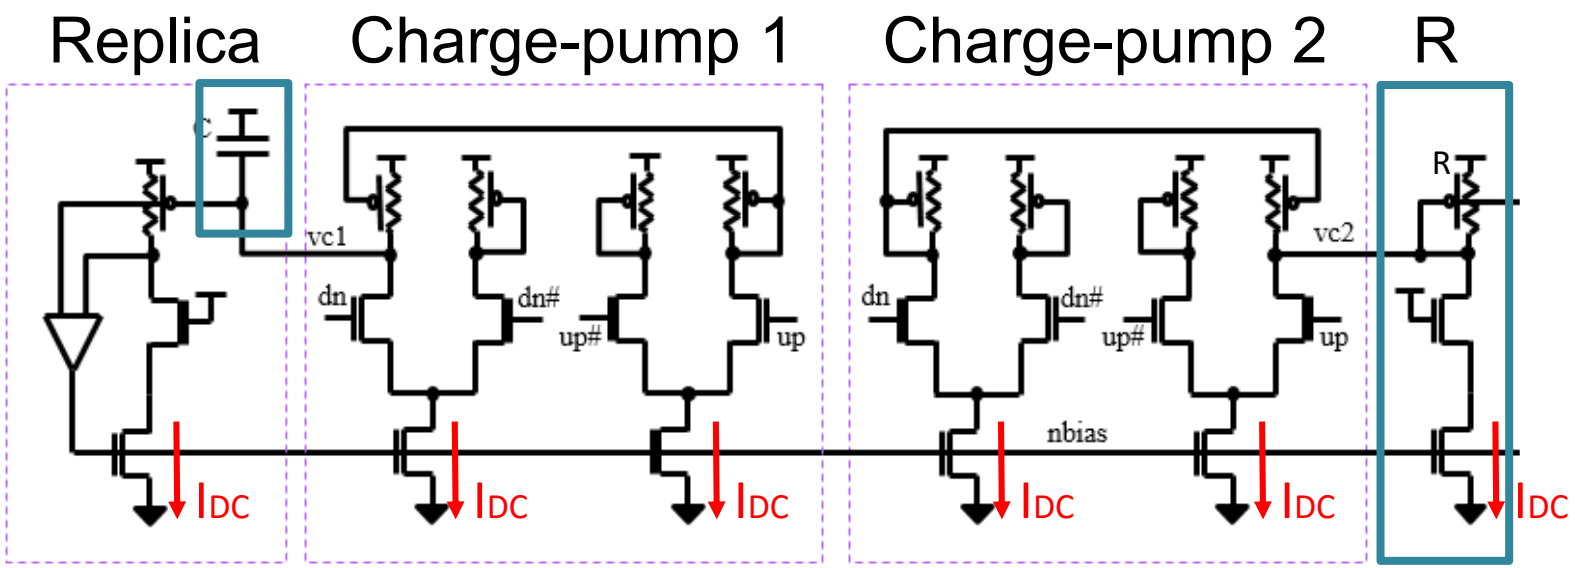

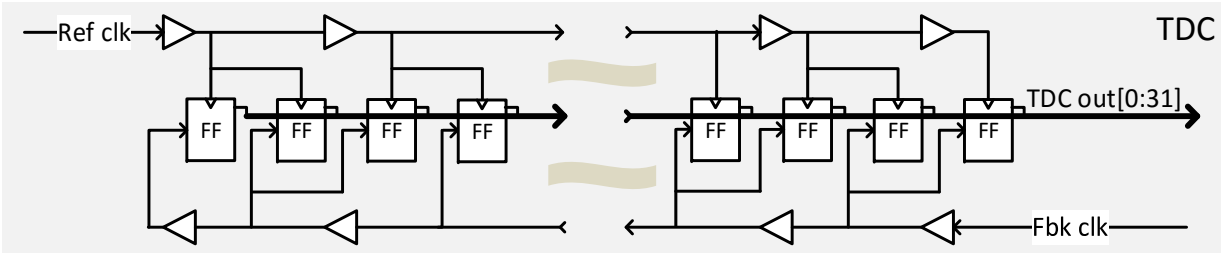

TDC

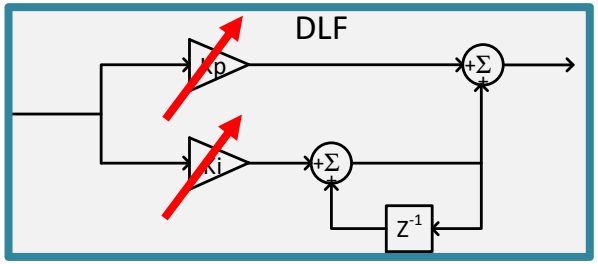

DLF

Display clocking using Fn DFG

- Display DFG is an LC based Fn Digital Frequency Generator (DFG)
  - Support DP, eDP 1.3/1.4 & HDMI frequencies
- 1 REF clock is used for SSC & non-SSC clock domains
  - SSC is generated within the DFG
- The DFG 1-ppm accuracy allows to generate any transfer-rate required by the display resolution and protocol
  - enabling usage of less I/O lanes thus increasing the power efficiency of the display block

Digital Divider-Less PLL TDC power

- In a Digital Divider-Less PLL architecture the DCO high frequency clock is driving directly the TDC circuit

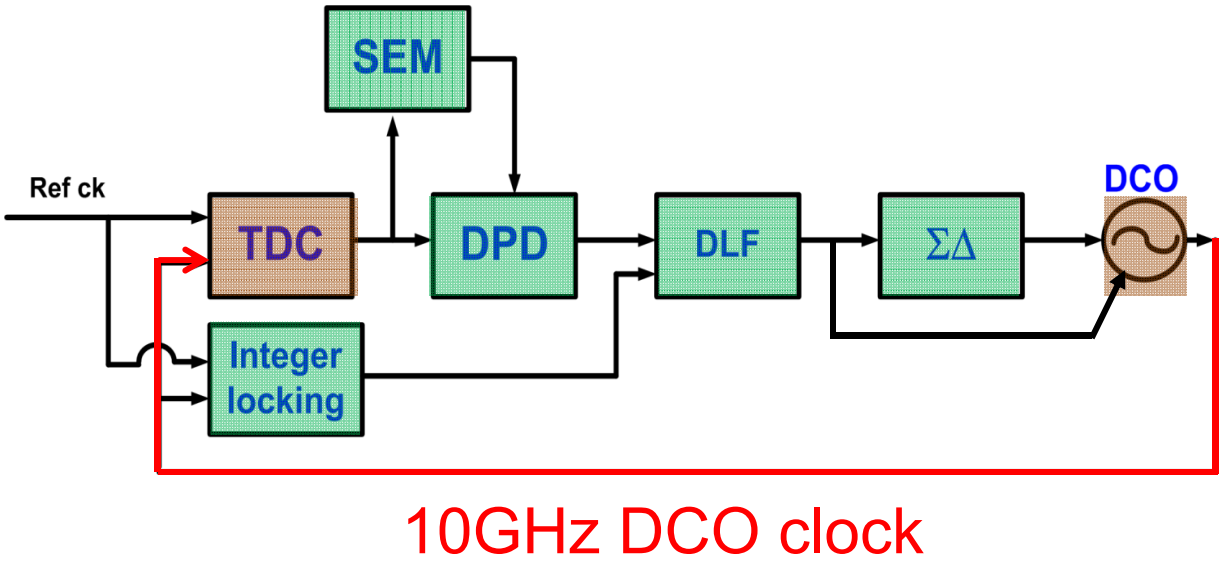

Digital Divider-Less PLL TDC power

- In a Digital Divider-Less PLL architecture the DCO high frequency clock is driving directly the TDC circuit
- The high-frequency toggling of the TDC delay line takes a significant part of the total DFG power

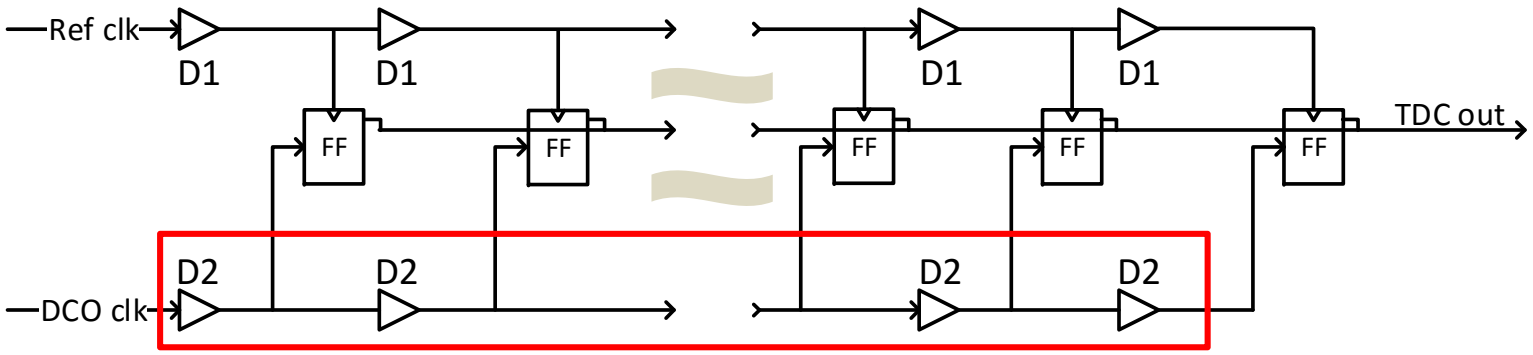

Vernier TDC

TDC gating technique to reduce DFG power

- Gating signal enables the DCO clock +/- 5% of the reference clock period surrounding the falling edge
  - 10X TDC power reduction

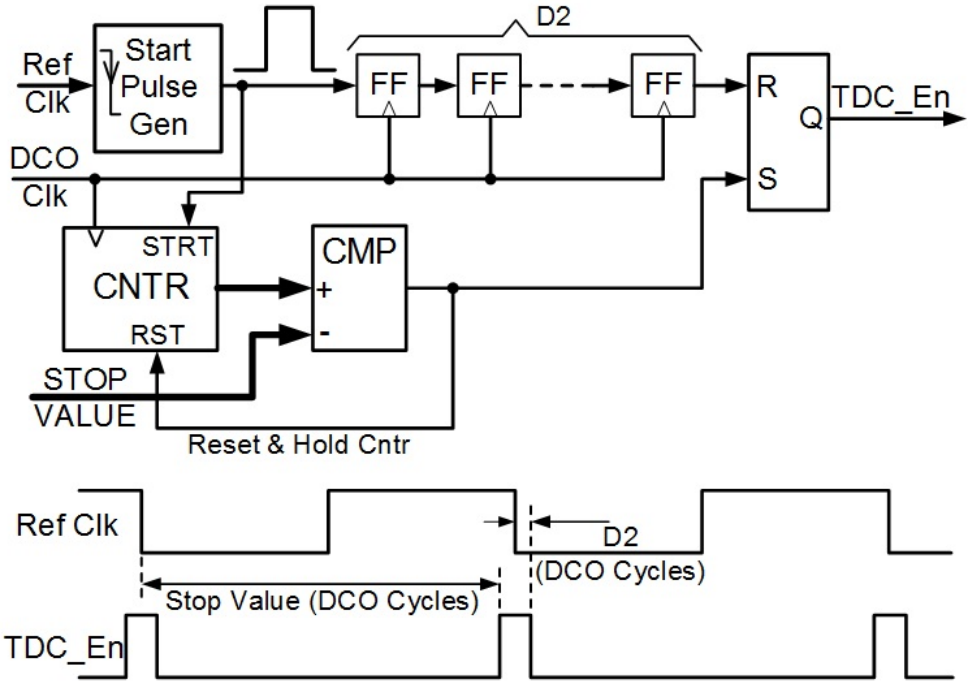

TDC gating technique to reduce DFG power

- Gating signal enables the DCO clock +/- 5% of the reference clock period surrounding the falling edge
  - 10X TDC power reduction

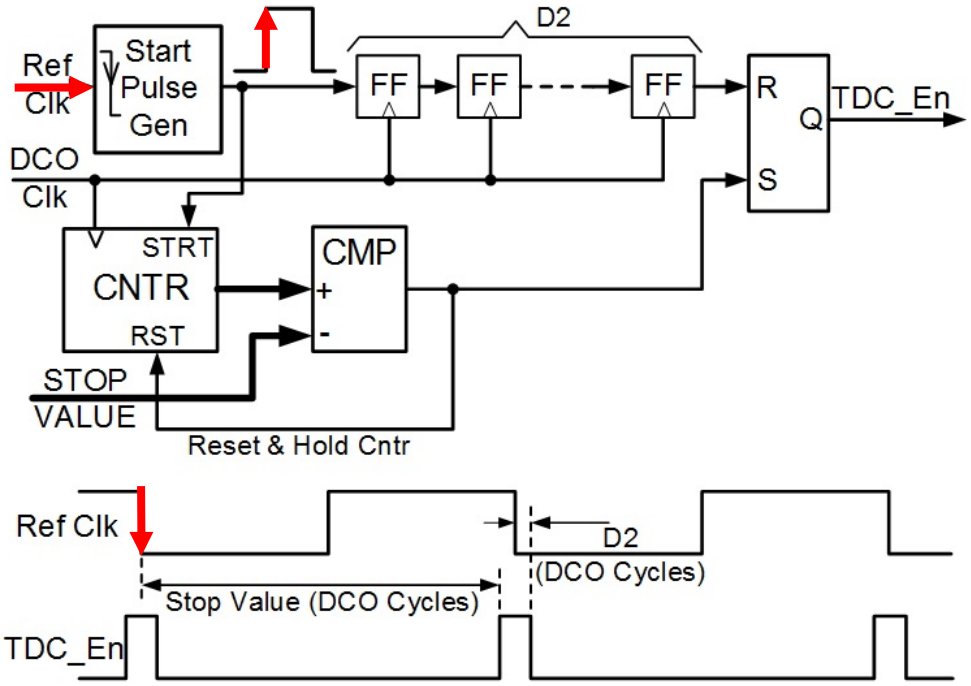

TDC gating technique to reduce DFG power

- Gating signal enables the DCO clock +/- 5% of the reference clock period surrounding the falling edge
  - 10X TDC power reduction

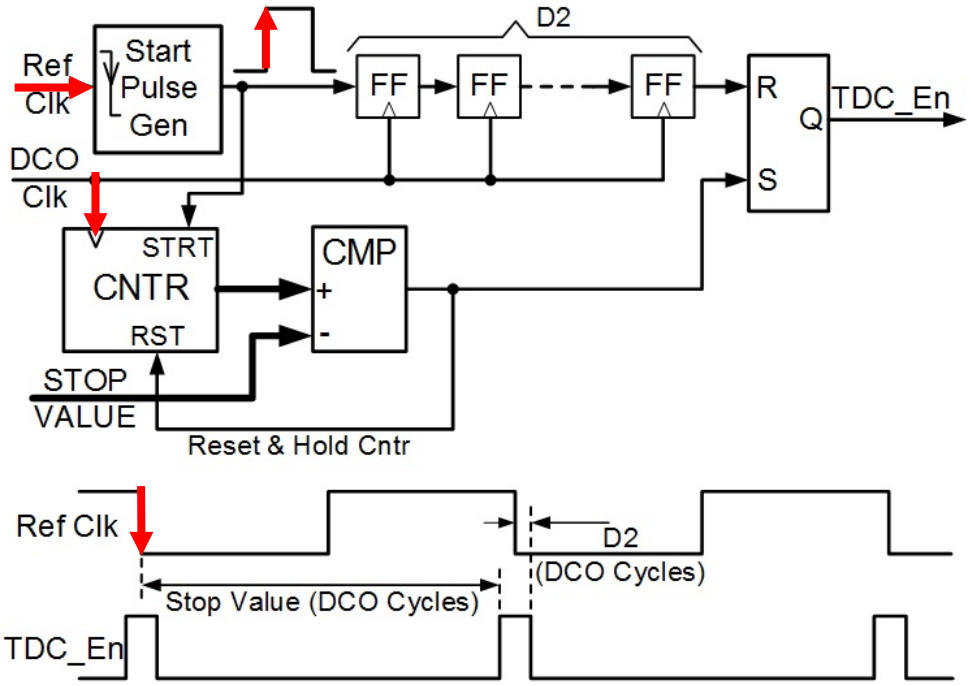

TDC gating technique to reduce DFG power

- Gating signal enables the DCO clock +/- 5% of the reference clock period surrounding the falling edge
  - 10X TDC power reduction

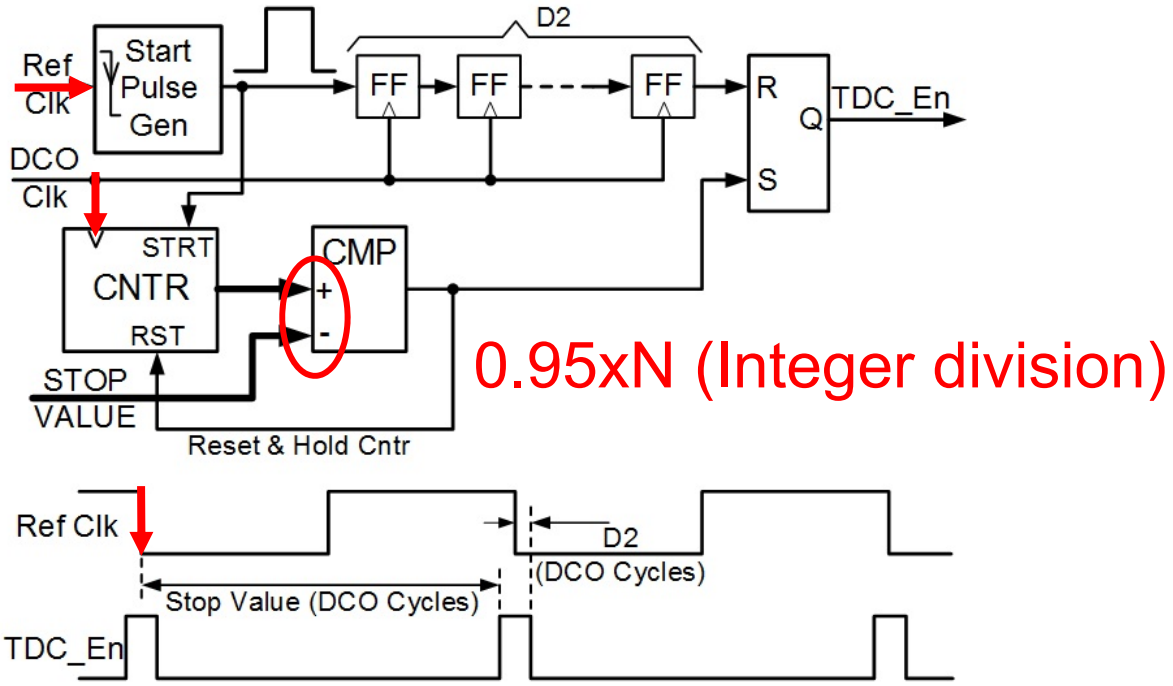

TDC gating technique to reduce DFG power

- Gating signal enables the DCO clock +/- 5% of the reference clock period surrounding the falling edge
  - 10X TDC power reduction

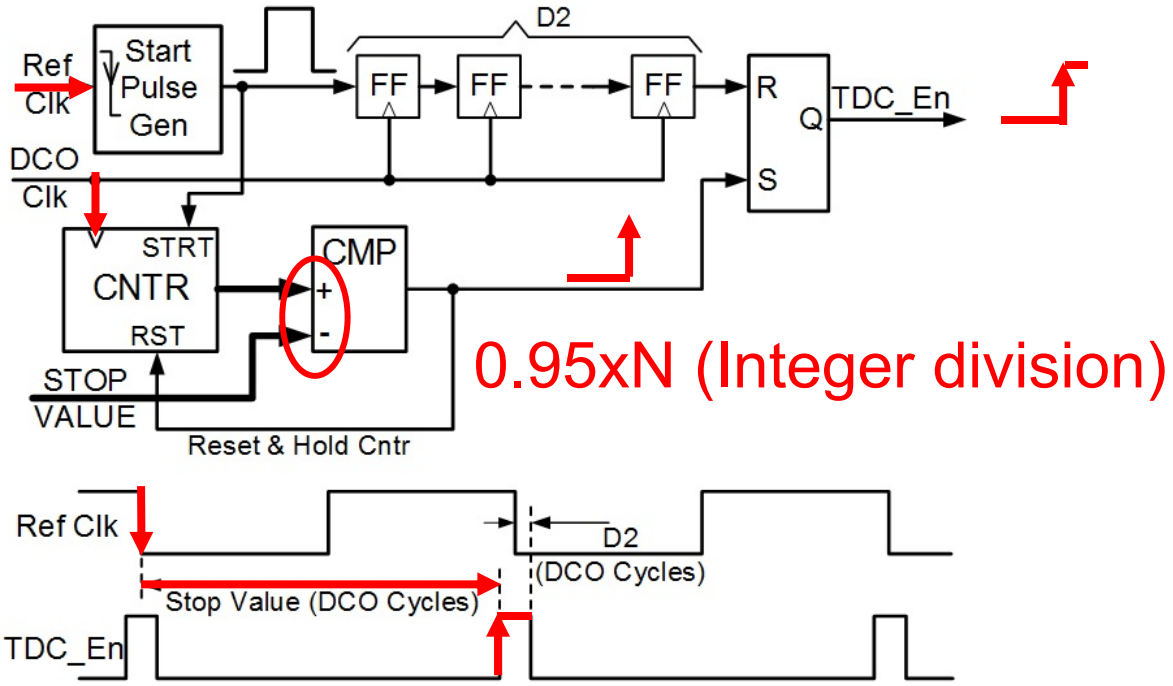

TDC gating technique to reduce DFG power

- Gating signal enables the DCO clock +/- 5% of the reference clock period surrounding the falling edge
  - 10X TDC power reduction

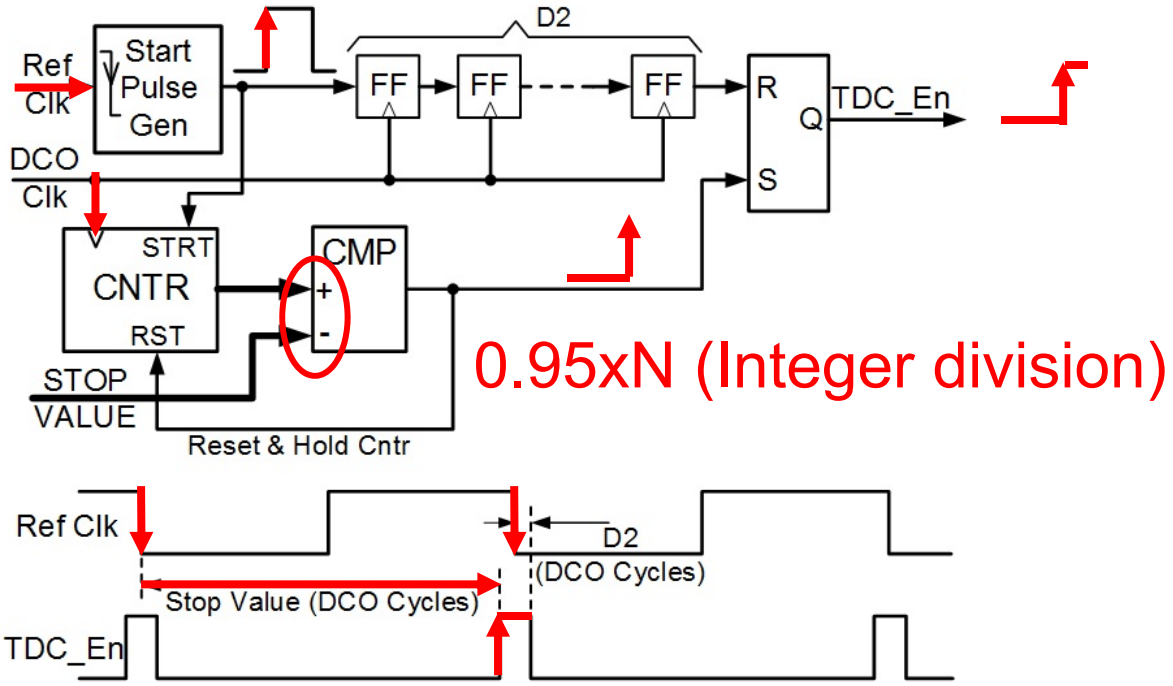

TDC gating technique to reduce DFG power

- Gating signal enables the DCO clock +/- 5% of the reference clock period surrounding the falling edge
  - 10X TDC power reduction

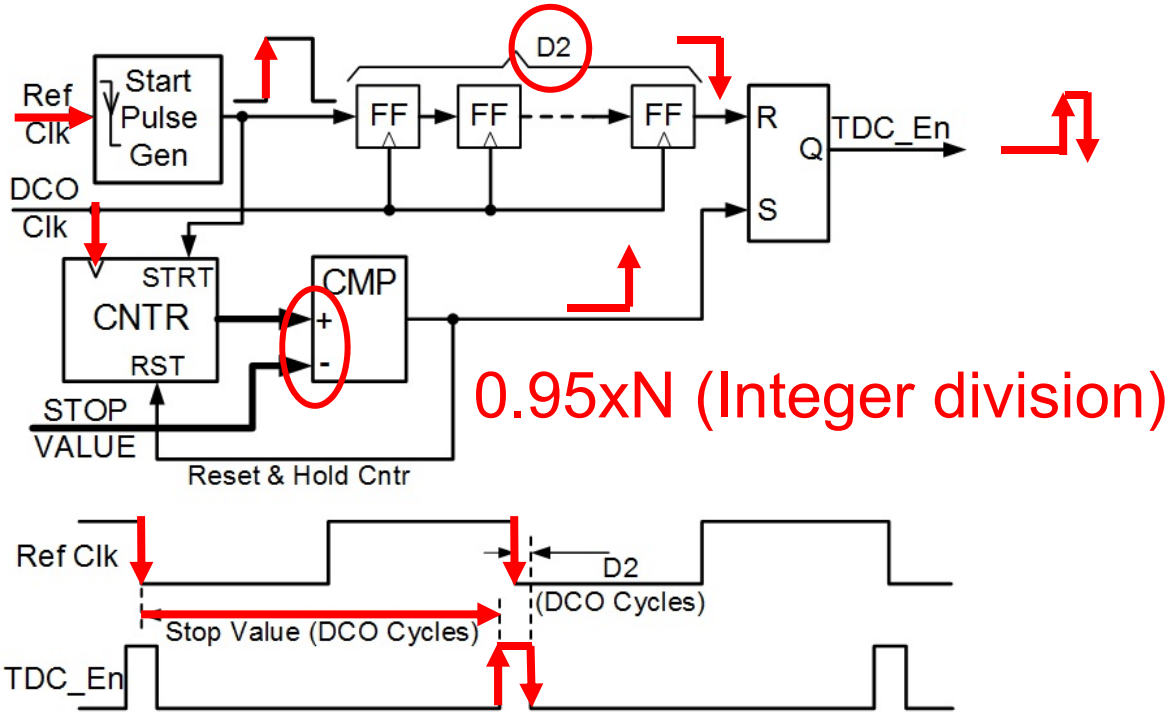

Island based Clock distribution NW

- Clock island approach enables multiple gated clock derivatives for power reduction
  - Support 58 different clock domains
  - Reduce the IA core total Cdyn by 2.3%

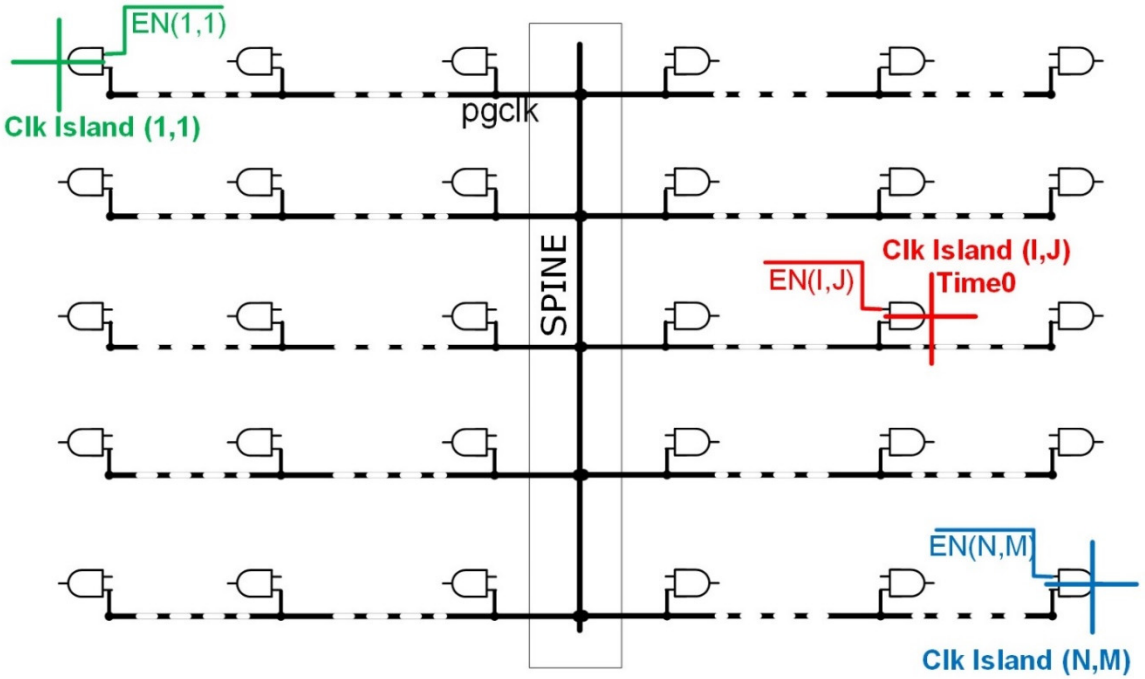

Island based Clock distribution NW

- Clock island approach enables multiple gated clock derivatives for power reduction
  - Support 58 different clock domains
  - Reduce the IA core total Cdyn by 2.3%

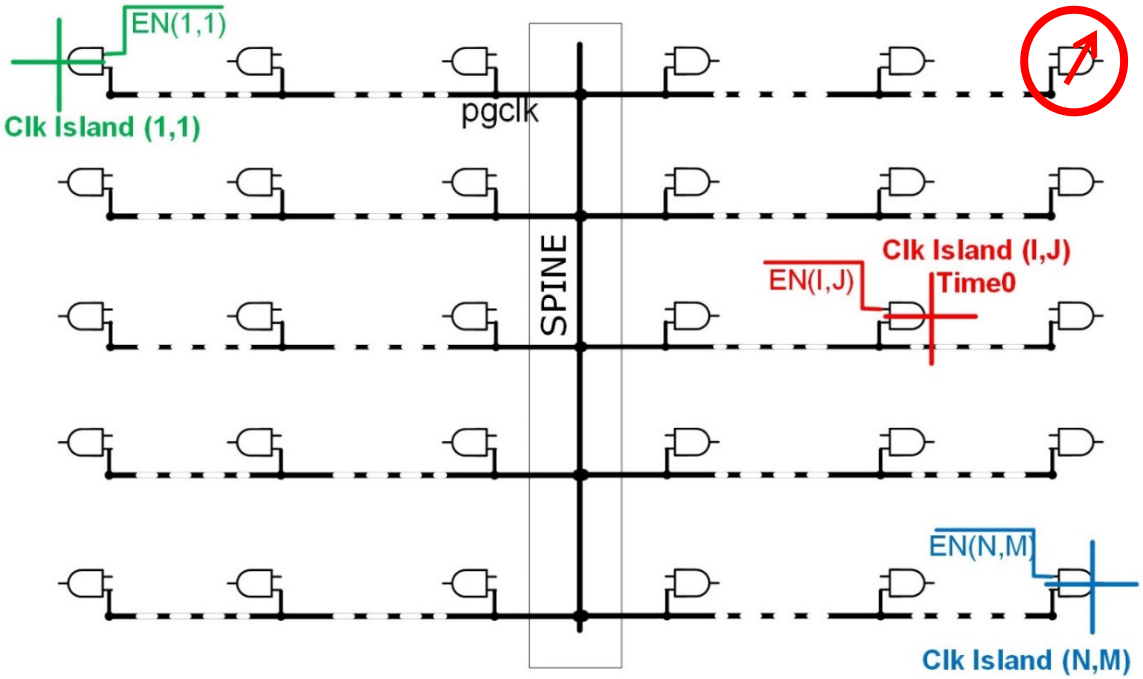

Tunable GB for power optimization & Global clock timing consolidation

SA clock distribution – MC & DDR

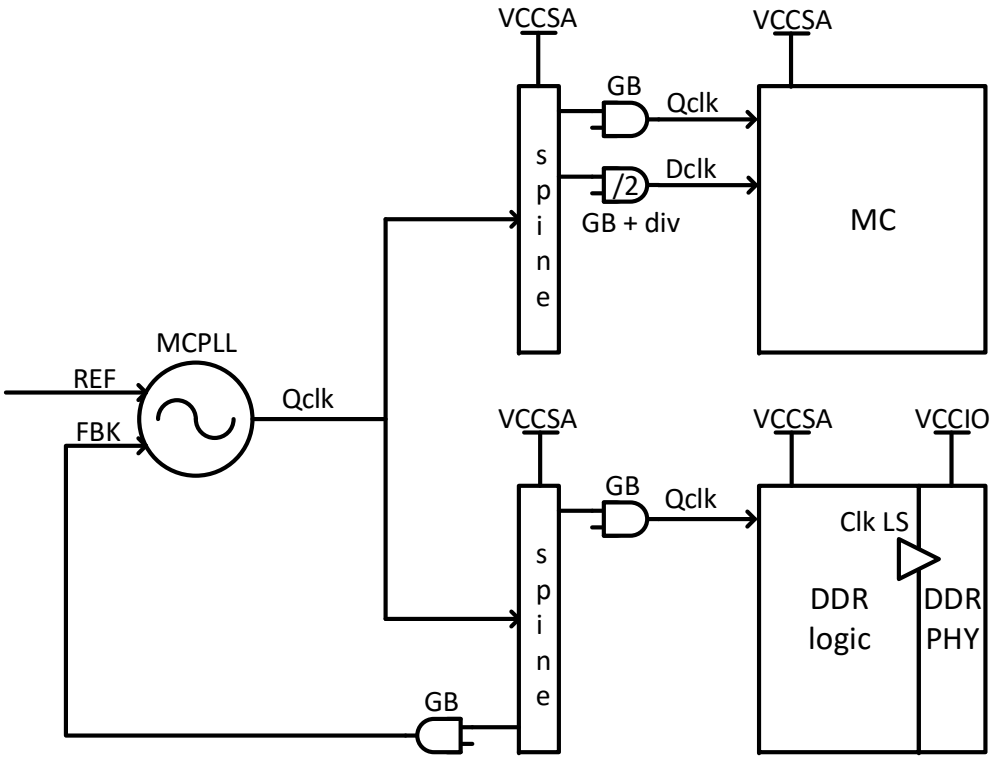

SA clock distribution – MC & DDR

Qclk, Dclk & FBK clock are aligned in time

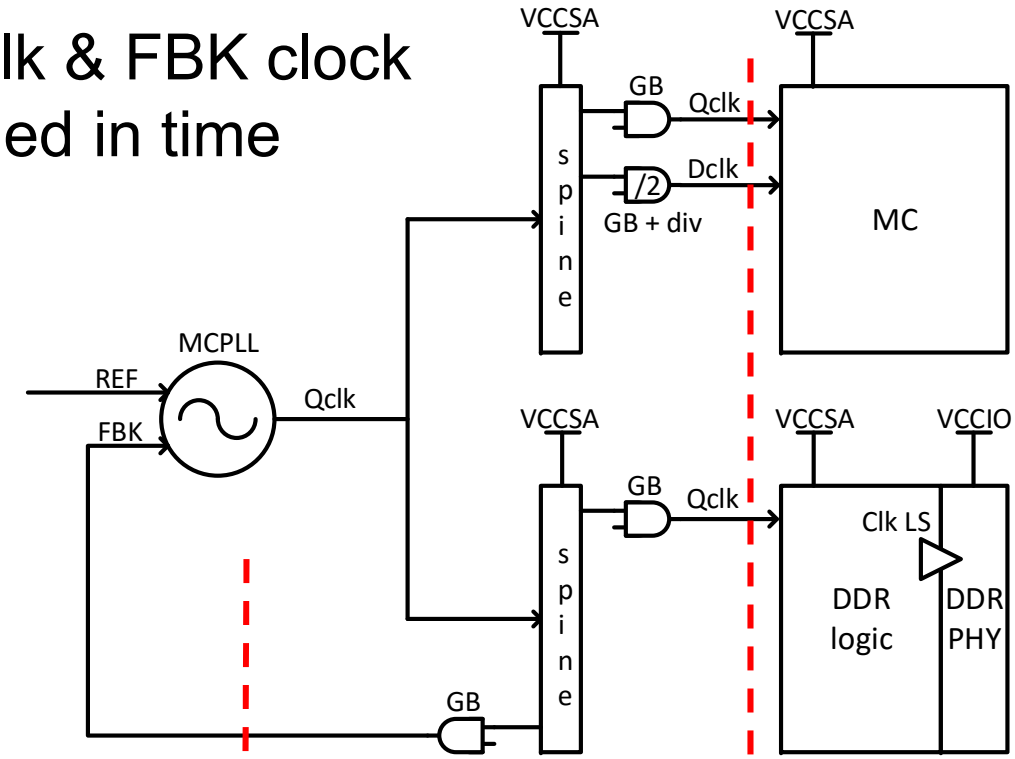

SA clock distribution - OPIO

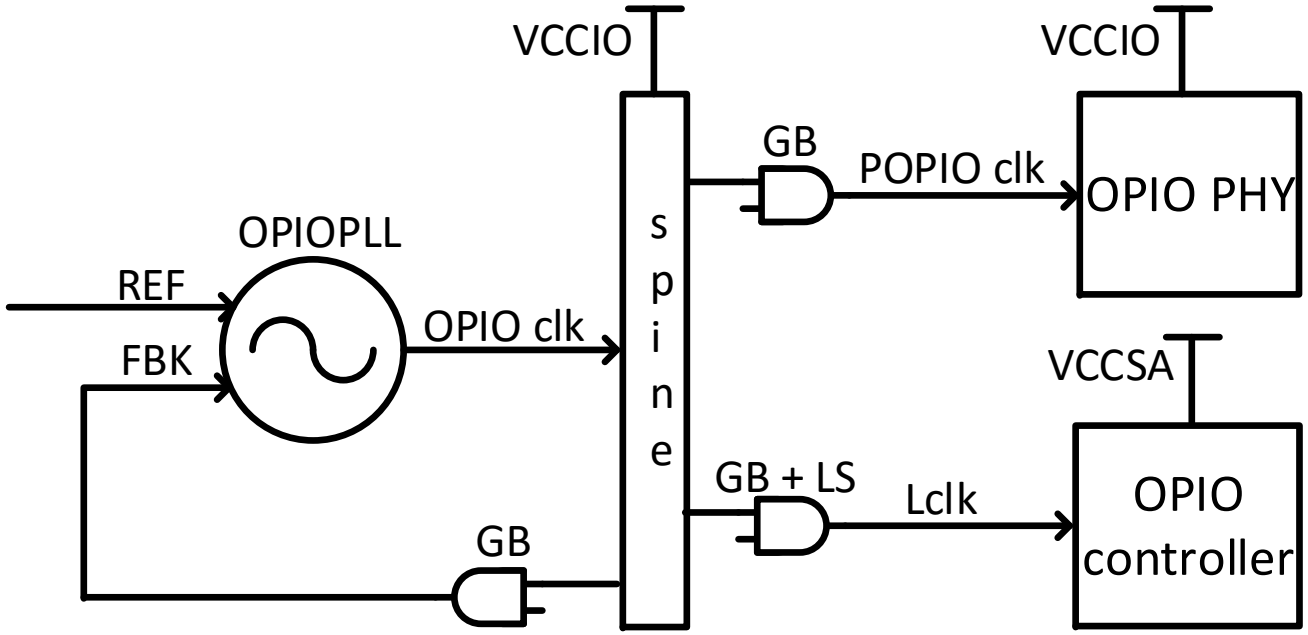

SA clock distribution - OPIO

POPIO clk, Lclk & FBK clk
are aligned in time

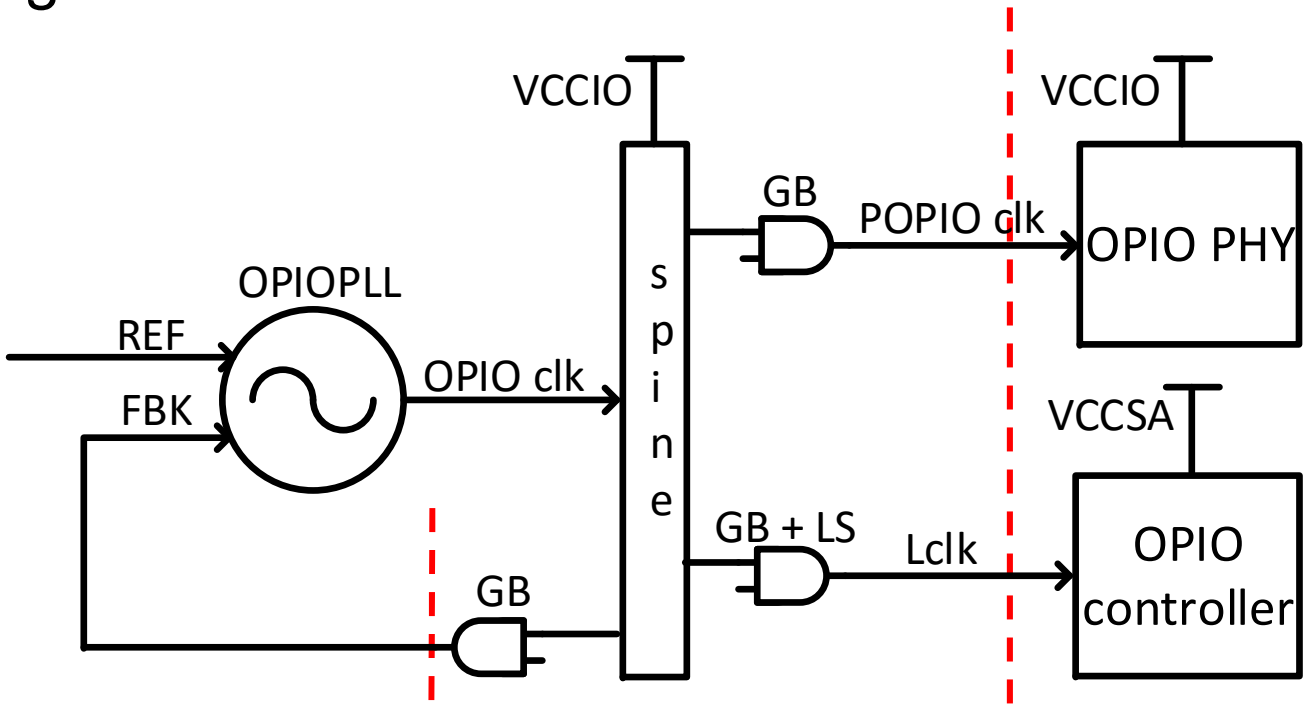

DDR Tx driver optimization

- DDR buffer-R_{ON} tuning dynamic range was optimized per die SKU
- Short DDR channel in U/Y SKU's allows reducing the driver strength by 33% while keeping full size for the DT die
- Driver strength reduction lowers C_{pad} which results in substantial switching power saving (for both write and read) and improves the read margin

$$\Delta P[W] = \Delta Cpad \times V^2 swing \times Freq \times K \times AF$$

$$K = number\ of\ \ pins$$

$$AF = Activity\ \ factor$$

DDR Match Current DLL power optimization

- DDR DLL is optimized for higher frequencies (not supporting ultra low frequencies <1067MHz)
- This allows cutting the DLL delay chain by half reducing the delay chain power by half
- To recover the 2X PI step 1 binary bit added to the PI circuit

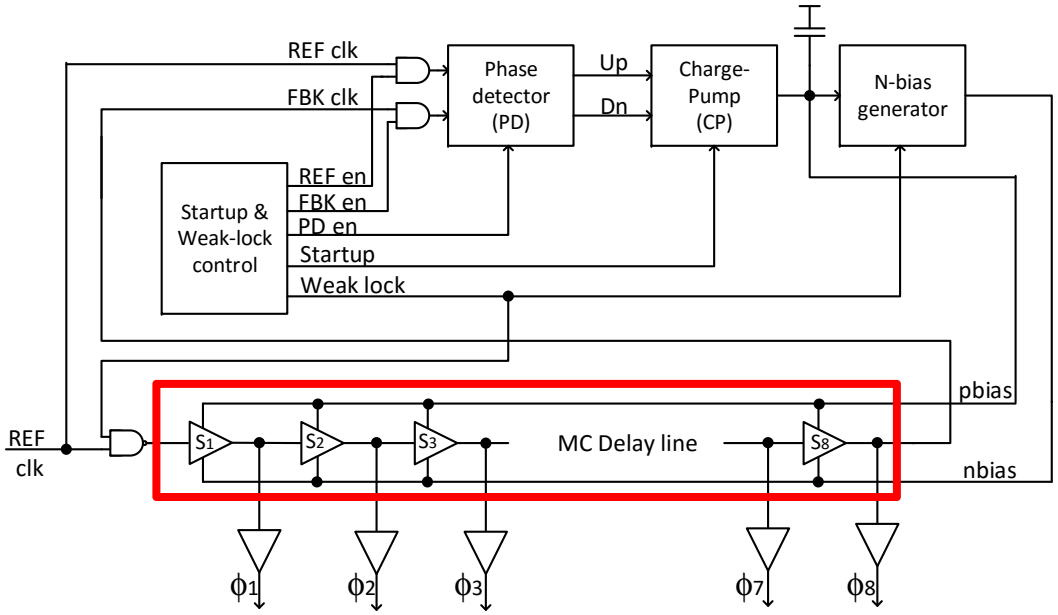

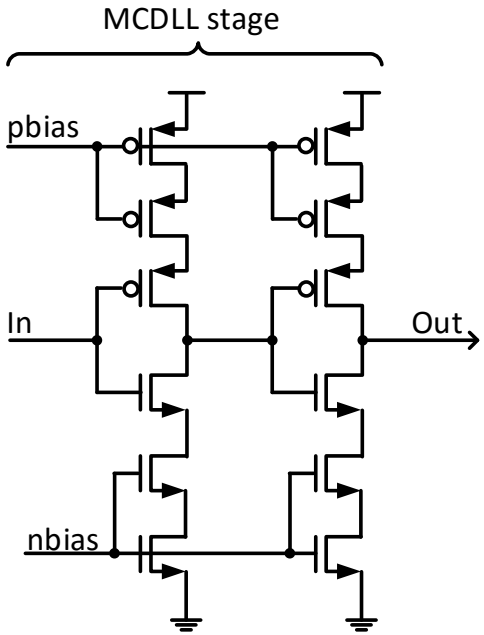

Dynamic gating of unused Phase Drivers

- Best power saving scenario: $\phi_{i,j} = \phi_{k,l} = \phi_{m,n}$
- No power saving scenario : $\phi_{i,j} = \phi_{k,l} = \phi_{m,n}$

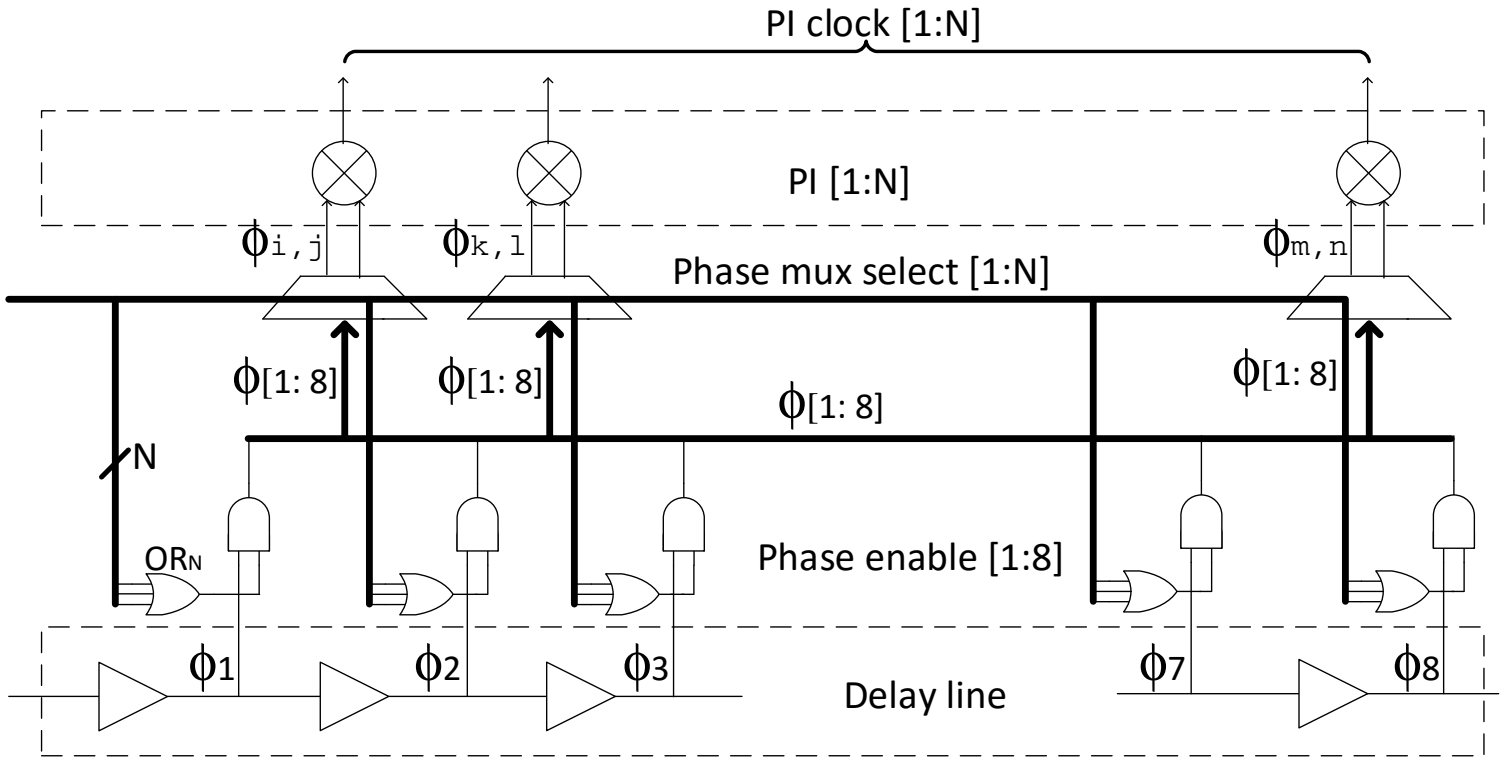

DLL Weak lock (WL) mode @ MC idle & no Tx/Rx

- REF clock en = 0 , FBK clock en = 0, CP is in OFF mode, Delay line is OFF

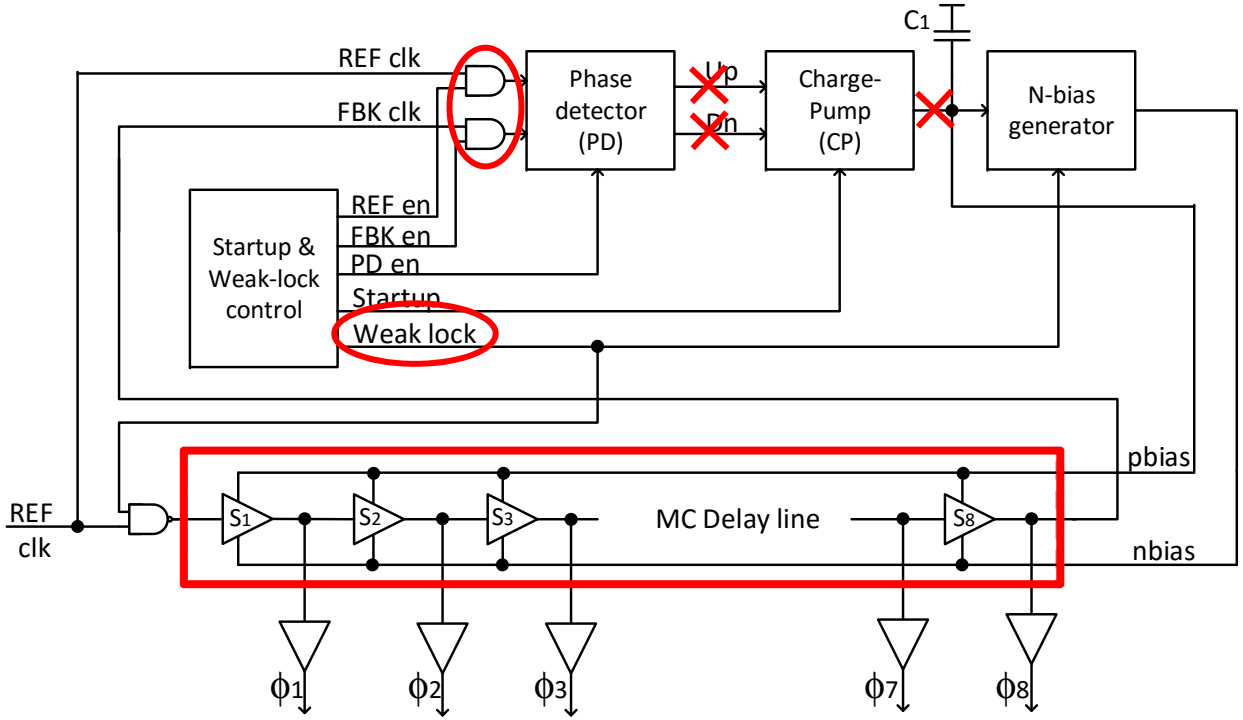

DLL Weak lock (WL) mode @ MC idle & no Tx/Rx

- REF clock en = 0 , FBK clock en = 0, CP is in OFF mode, Delay line is OFF
- Vcntl is stored in C1

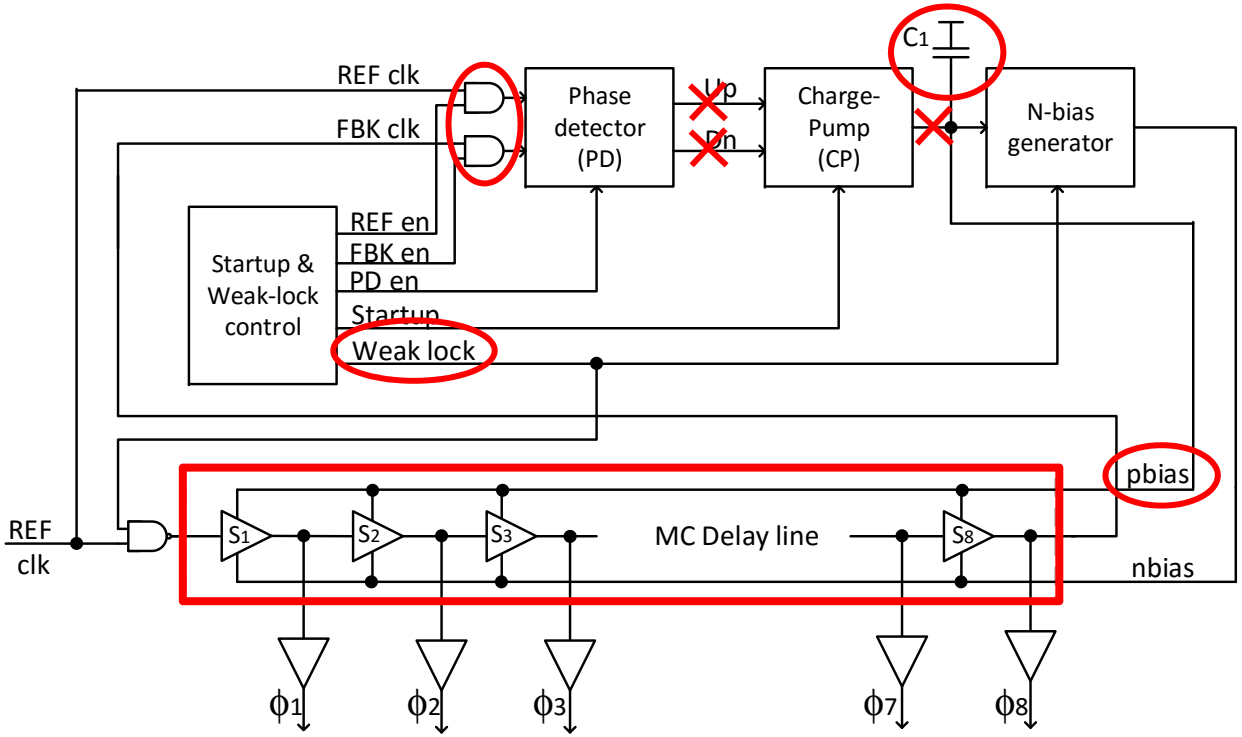

DLL Weak lock (WL) mode @ MC idle & no Tx/Rx

- REF clock en = 0 , FBK clock en = 0, CP is in OFF mode, Delay line is OFF
- Vcntl is stored in C1
- N-bias is generated by half of the current replica

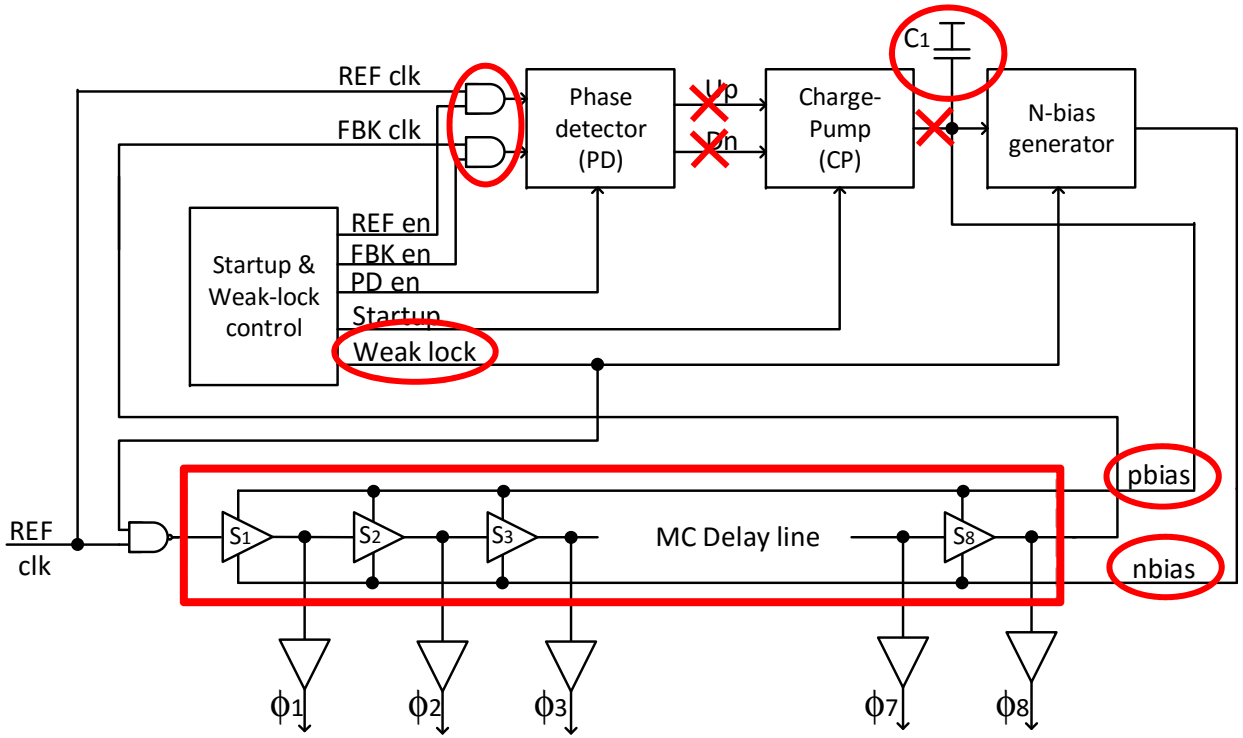

DLL Weak lock (WL) mode @ MC idle & no Tx/Rx

- REF clock en = 0 , FBK clock en = 0, CP is in OFF mode, Delay line is OFF
- Vcntl is stored in C1
- N-bias is generated by half of the current replica
- To maintain WL, WL activity factor is configured to 40%

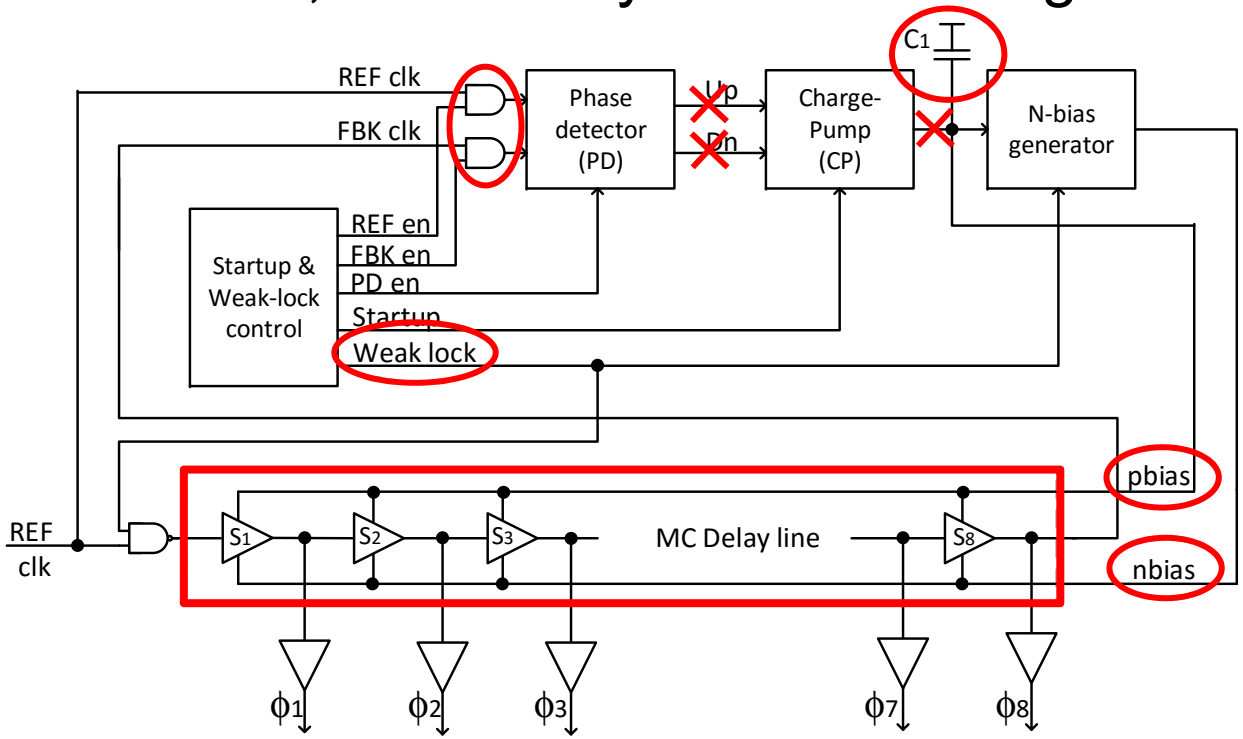

Video playback power optimization

- Compute architecture of Intel® processor graphics gen9 is based on a single slice with three sub-slices, for a total of 24 EUs.

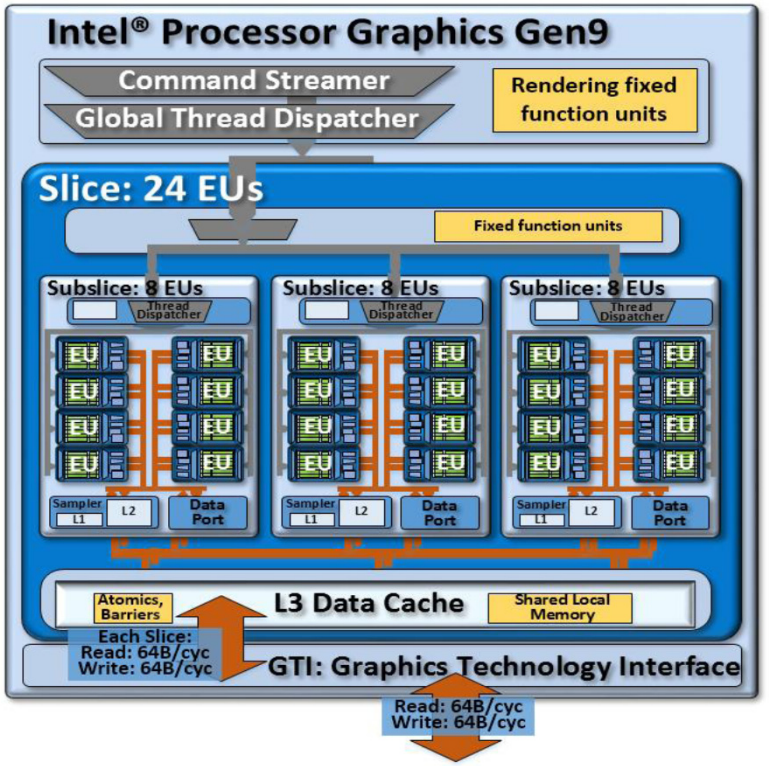

Video playback power optimization

- Compute architecture of Intel® processor graphics gen9 is based on a single slice with three sub-slices, for a total of 24 EUs.
- Heavy parts of the VPB execution were moved to dedicated blocks that were optimize for low power & memory efficiency

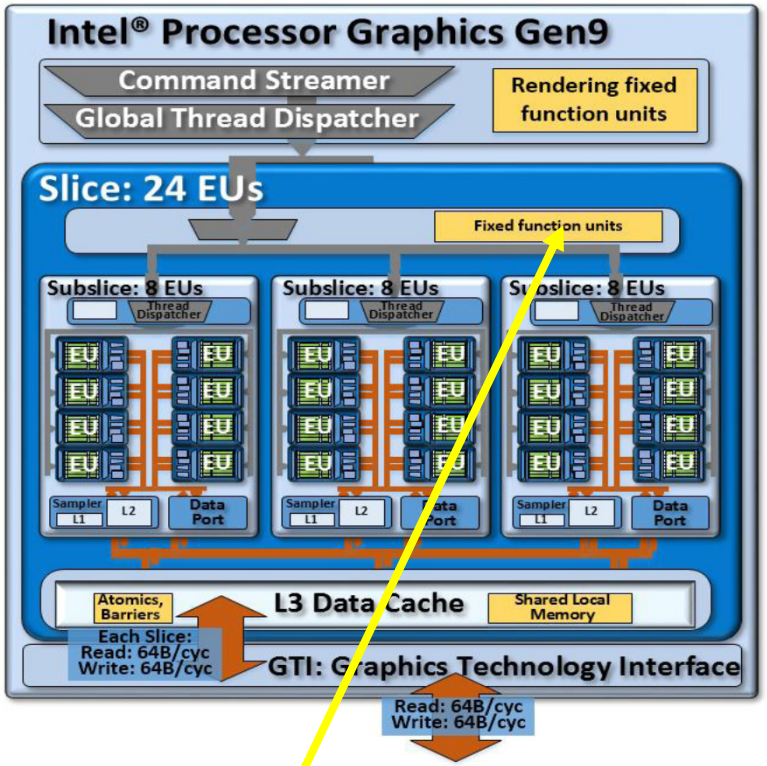

Fixed Function Unit (Media block)

Video playback power optimization

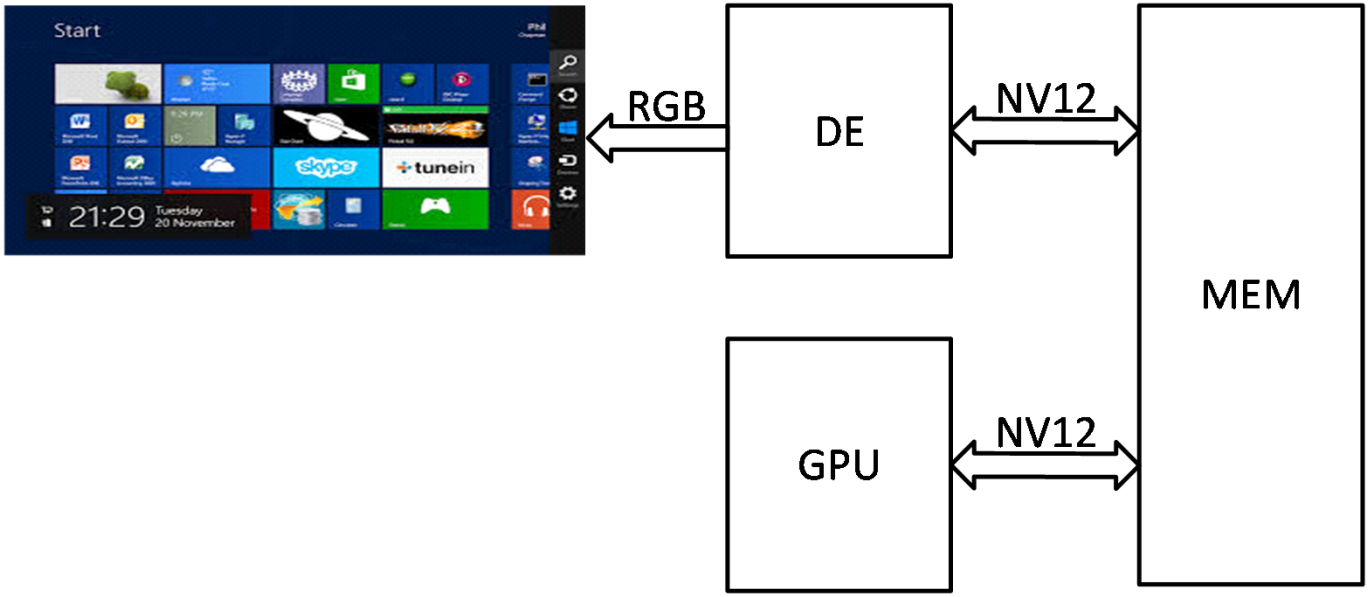

Video playback power optimization

- Use Full-HD data in NV12 format for all memory transactions

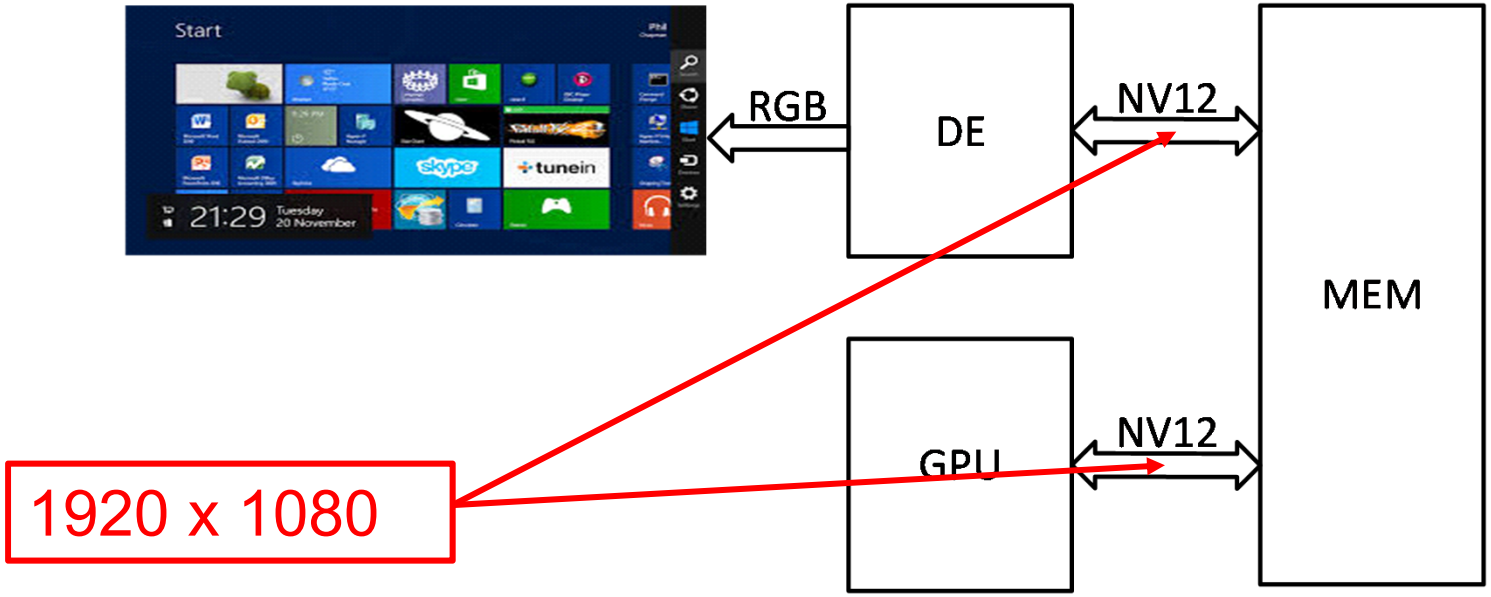

Video playback power optimization

- Use Full-HD data in NV12 format for all memory transactions
  - NV12 to RGB conversion is done in the DE

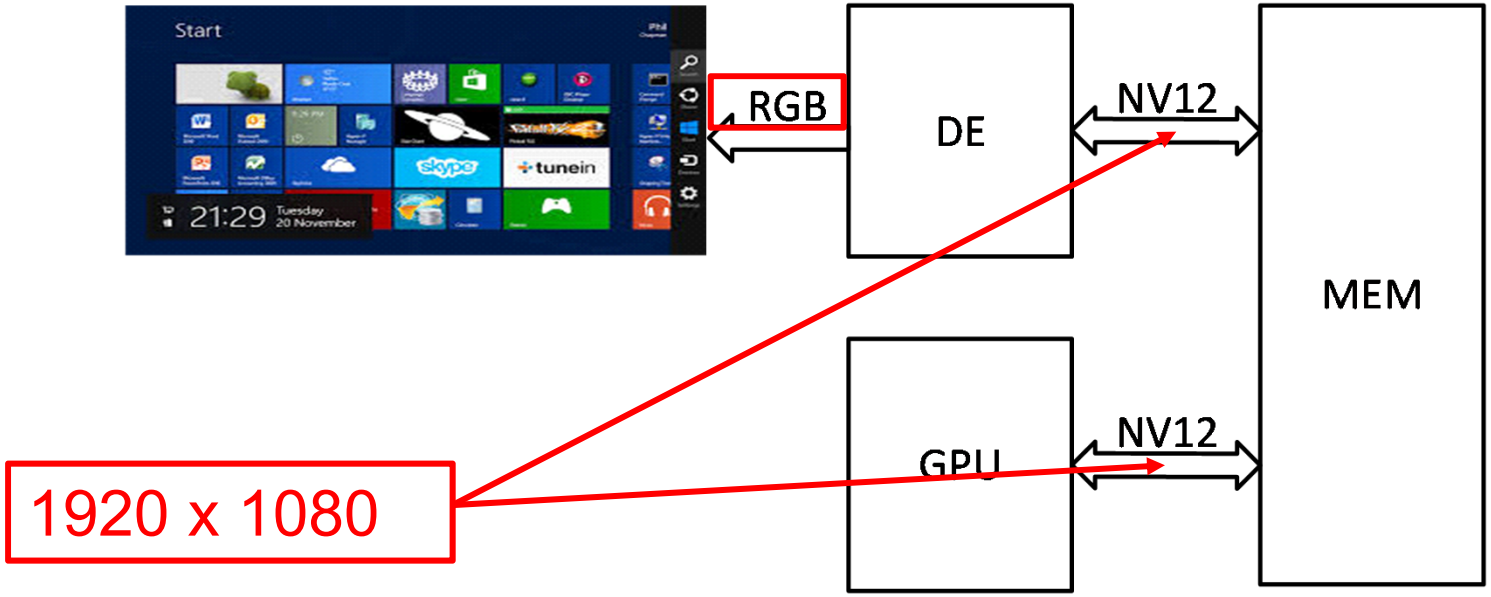

Video playback power optimization

- Use Full-HD data in NV12 format for all memory transactions
  - NV12 to RGB conversion is done in the DE
- Image resolution is scaled up at the end of the Display pipe

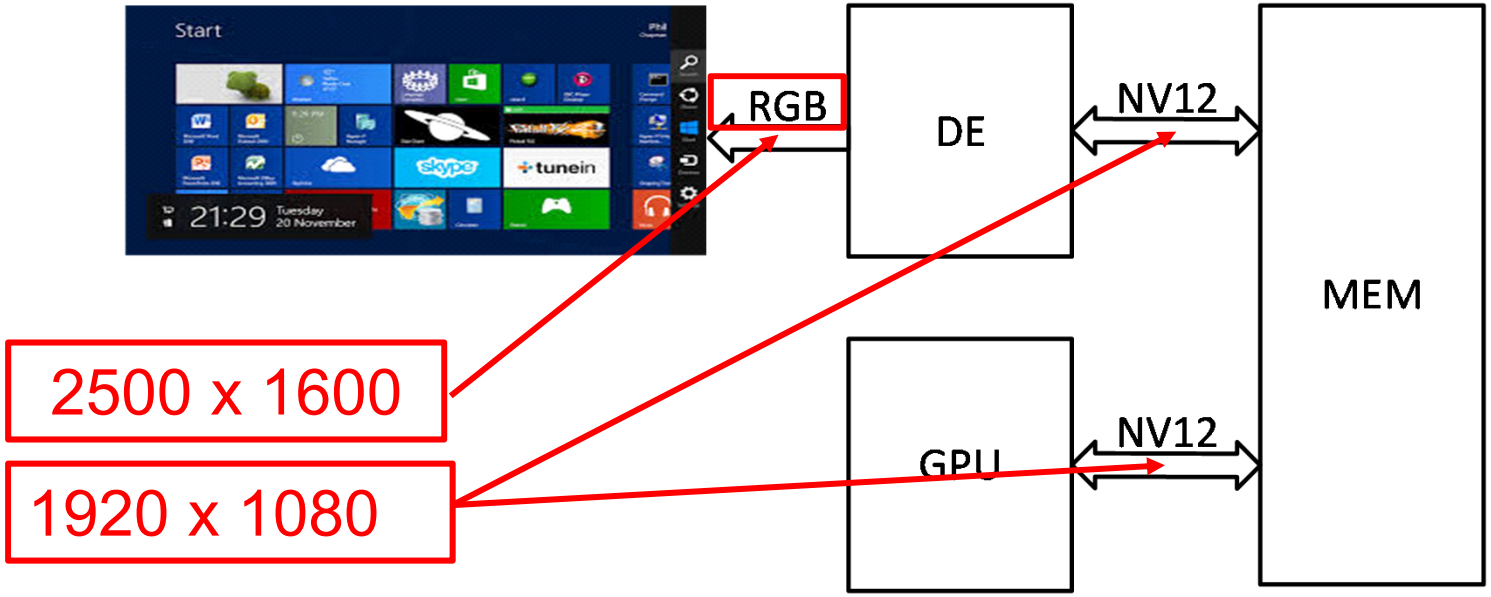

Video playback power optimization

- Use Full-HD data in NV12 format for all memory transactions
  - NV12 to RGB conversion is done in the DE
- Image resolution is scaled up at the end of the Display pipe
- Optimized DE buffer size to reduce wakeup rates for 2K & 4K panels

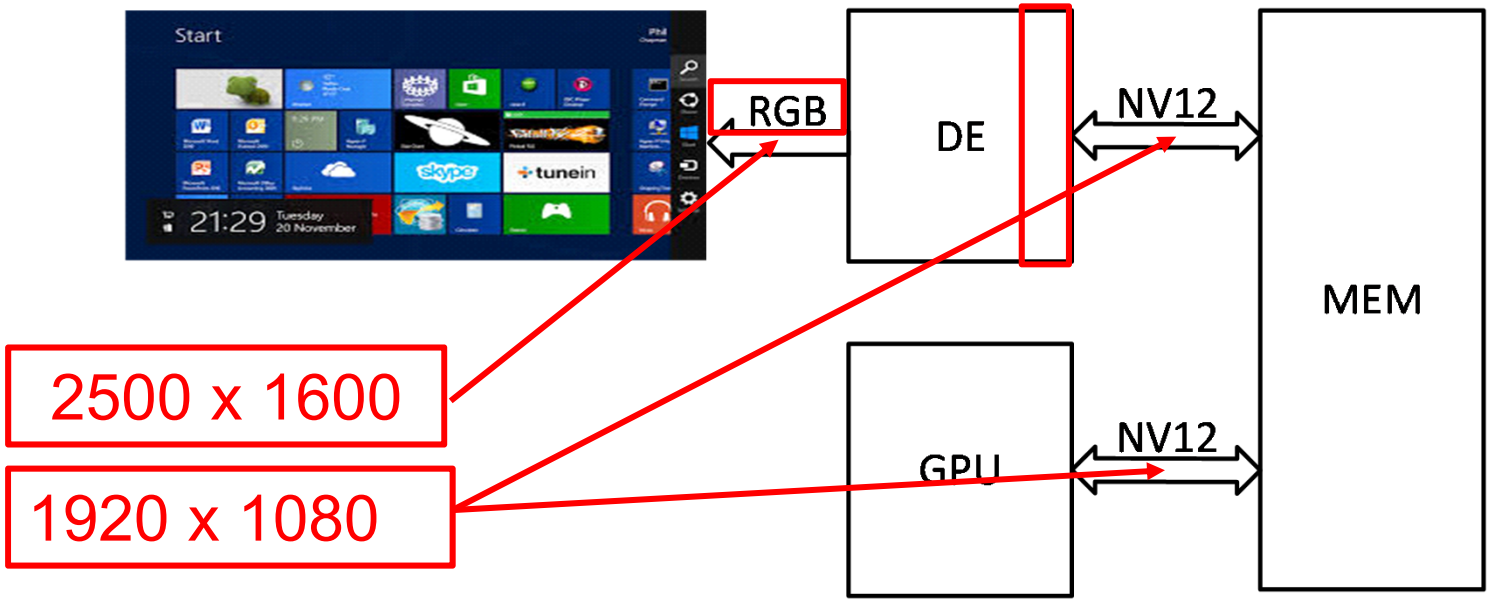

Video playback power optimization

- Use Full-HD data in NV12 format for all memory transactions
  - NV12 to RGB conversion is done in the DE
- Image resolution is scaled up at the end of the Display pipe
- Optimized DE buffer size to reduce wakeup rates for 2K & 4K panels

Memory traffic reduced by 5X

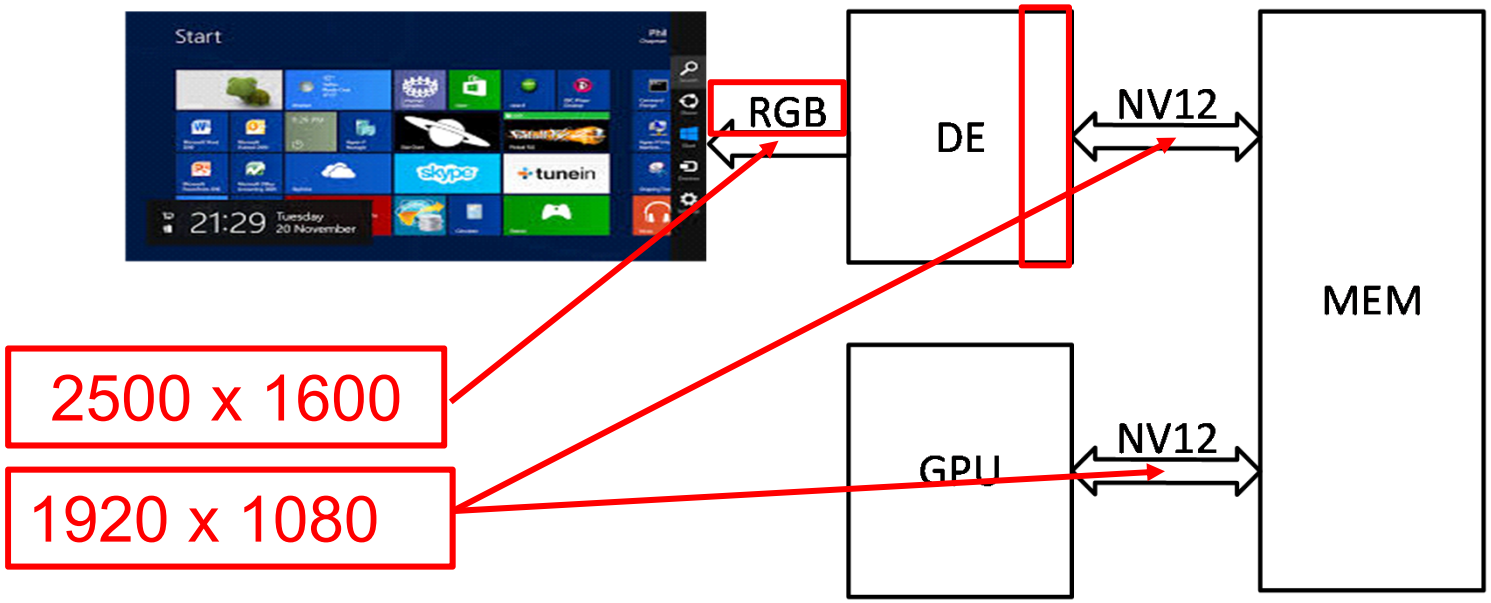

Memory sub-system power optimization

- Enabling Dynamic Voltage & Frequency Scaling (DVFS) for the System Agent (SA)
- Reduce SA & IO voltages
- Enable C9/C10 with Panel Self Refresh (PSR)
  - Reduce latencies for Save/Restore flows

Post Si power vs. 14nm previous generation

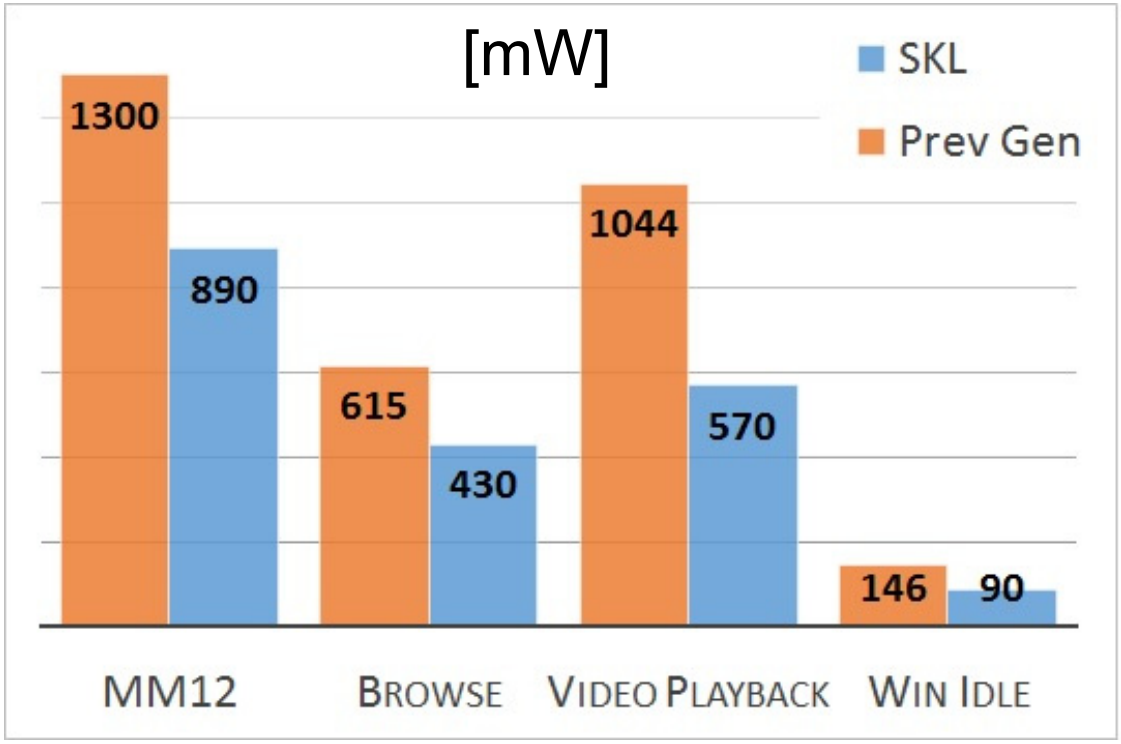

Post Si power vs. 14nm previous generation

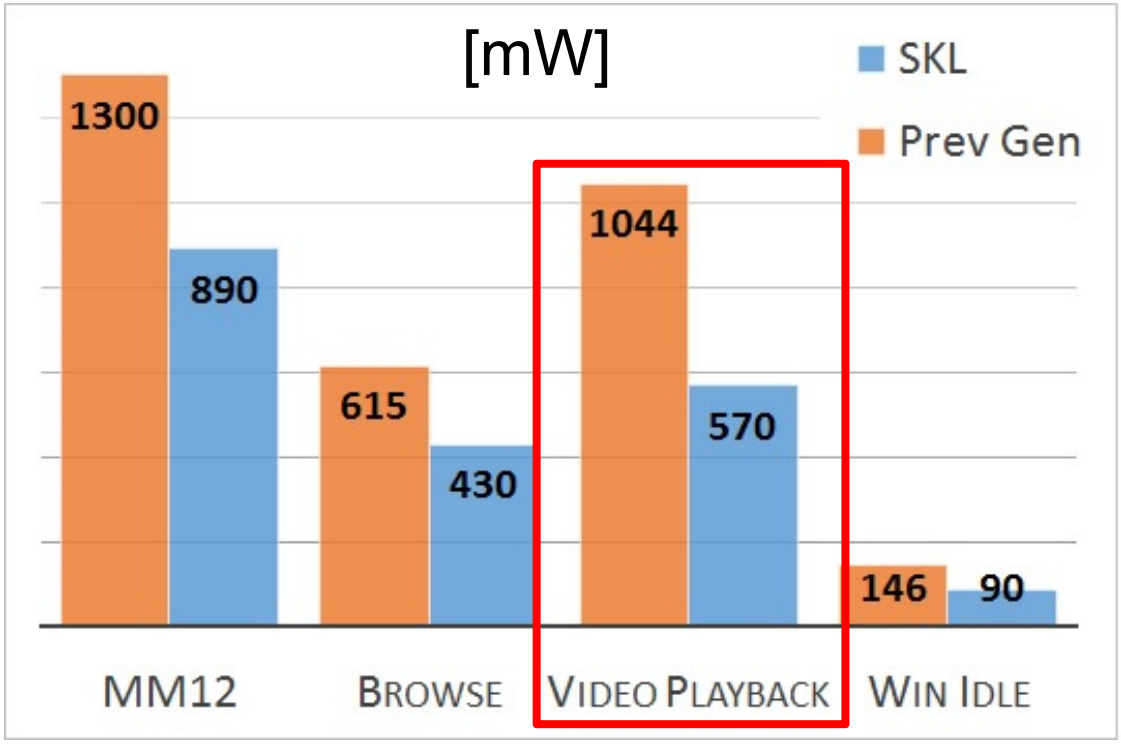

- VPB power reduction is 45% vs. 14nm previous generation
  - 21% by general HW optimization
  - 24% by specific VPB optimization

Summary

- Intel's 6th generation core ™ has considerable improved performance and power dissipation using innovative techniques
  - Digital PLLs & low-power clocking
  - Aggressive power gating
  - Low power I/O PHY's
  - Specific VPB optimization
- The chip floorplan provides various scalability alternatives to increase the range of price-performance-power optimization
- The part is at high volume manufacturing using Intel's 14nm process

Increasing the Performance of a 28nm x86-64 Microprocessor Through System Power Management

Aaron Grenat, Sriram Sundaram, Stephen Kosonocky, Ravinder Rachala, Sriram Sambamurthy, Steven Liepe, Miguel Rodrigues, Thomas Burd, Adam Clark, Mike Austin, Samuel Naffziger

Outline

- Shadow P-states
- Reliability Tracker
- Digital LDO Regulation
- Skin Temperature Aware Power Management
- Boot-time power supply calibration
- Conclusion and measured Si performance results

Overview

- 28nm Bulk CMOS
- 12 metal layers (eight 1x, one 2x, one 4x, two 16x)
- Bristol Ridge (BR)
  - Area: 250.04 mm^2
  - Transistors: 3.1B
- XV Module (core pair)
  - Area: 14.48 mm^2
  - Transistors: 102M
- Project goals
  - Use Carrizo Masks
  - Provide 10-15% performance improvement

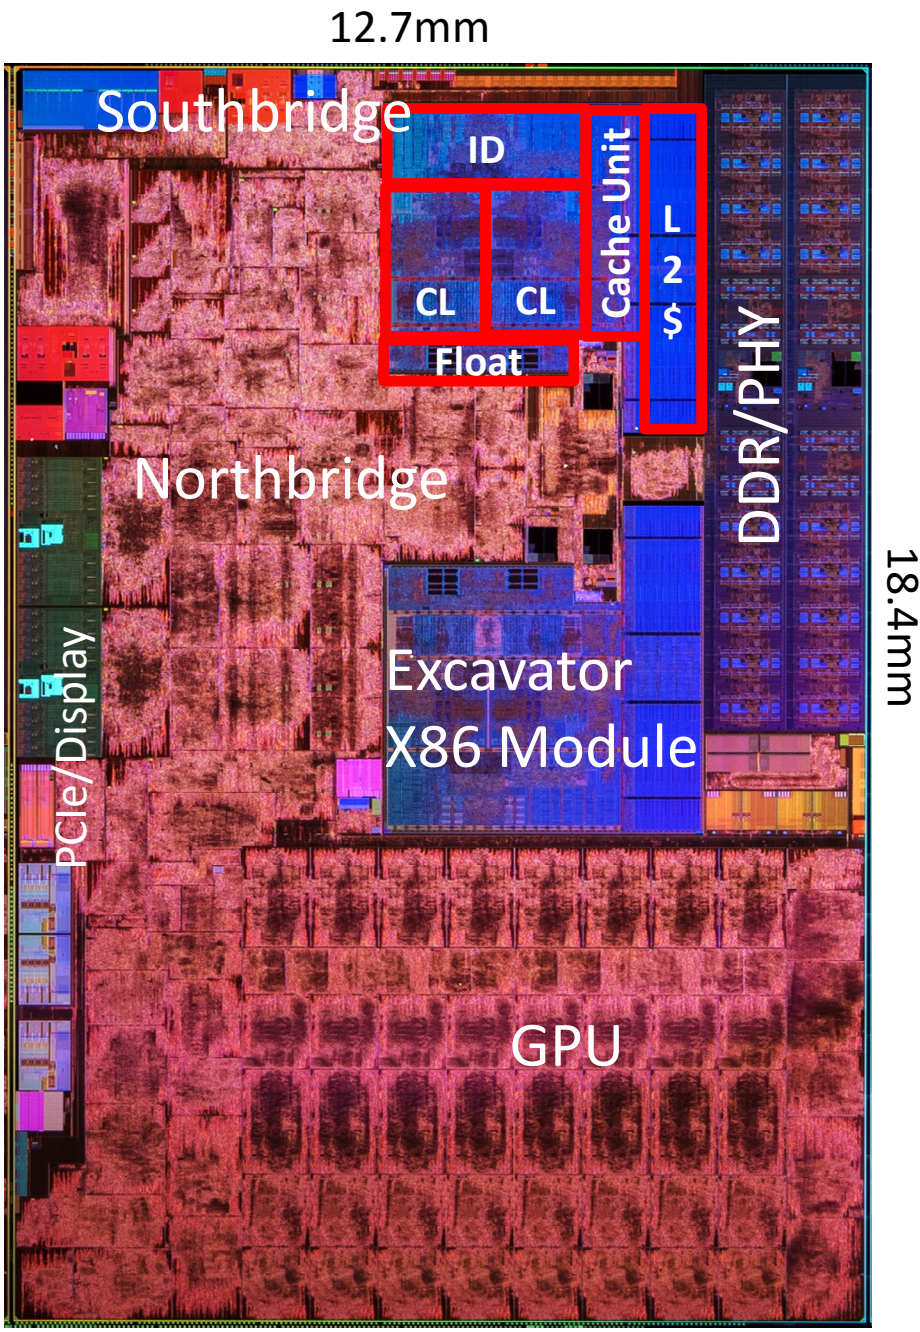

Bristol Ridge Major Structures

- Four Excavator cores
- Eight Radeon™ Graphics Core Next (GCN) Compute Cores
  - Support HSA compute acceleration
- High throughput multimedia accelerators
  - HEVC(H.265) support
  - re-architected for low power video playback
  - up to 3.5X improvement on 1080p video transcode
- Integrated Platform security processor
  - platform secure boot
- Integrated Southbridge
- High performance I/O connectivity
  - 128 bit DDR channel supports DDR3 and DDR4
  - 12 lanes of PCIe GEN3
  - 3 independent display heads capable of DP, DVI, HDMI
  - USB2, USB3, SATA, UART, I2S connectivity

Shadow Pstates

- Process variations that define the speed/power characteristics of the part, VRM tolerance and package characteristics dictate how much voltage every part on the system needs to achieve certain target frequencies

- AVFS gives us the part-specific voltage necessary to achieve target frequencies

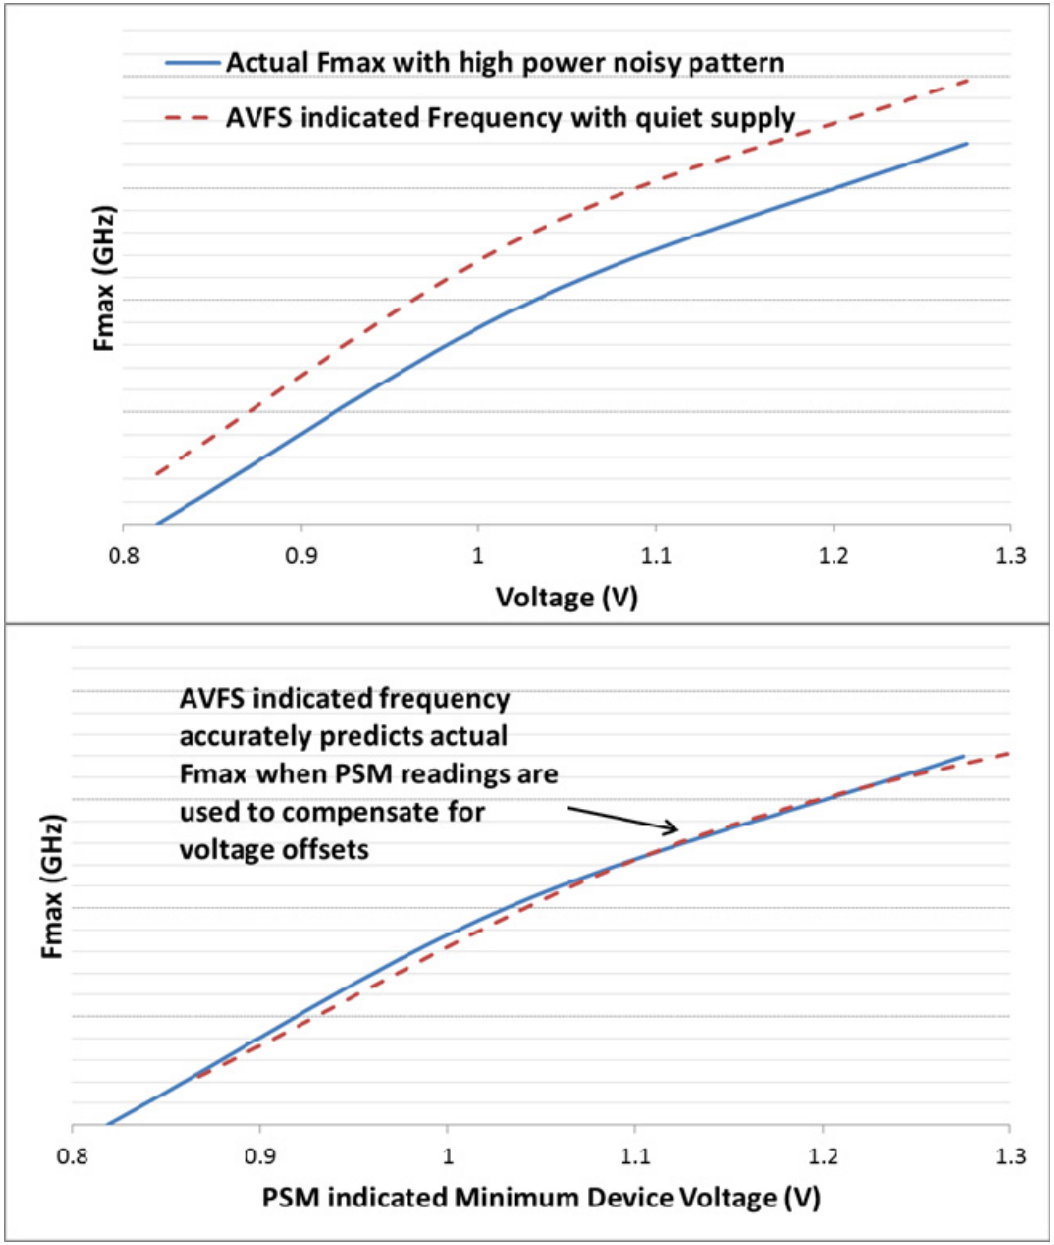

Shadow Pstates

- SMU firmware checks that the target frequency + xMHz (increased frequency) is achievable using AVFS and BTC's frequency-voltage curves, and that other system infrastructure (EDC, Vmax) limits will be met at this higher voltage/frequency.

- If the checks are satisfied for the platform, the frequency of the IP is increased by that x Mhz thereby gaining more performance.

- We obtained an additional 100 MHz peak boost frequency on Bristol, using shadow Pstates, when compared to traditional binning

Reliability Limiters

- Critical reliability failure modes with a strong function on operating conditions (Voltage, Frequency, Temperature):
  - Time-dependent Dielectric Breakdown (TDDB)
  - Electromigration (EM)

- Failure Rate expressed in FIT, sum of individual terms:
  - *1 FIT = 1 Failure in a billion operating hours.*
  - $FIT_{Total} = FIT_{TDDB} + FIT_{EM}$

- For a specific set of usage model, product lifetime, & FIT_{TOTAL} assumptions, it is additive in time for small FIT_{TDDB} (& FIT_{EM}):

$$FIT_{TDDB|Lifetime} = \sum_{i}^{Lifetime} FIT_{TDDBi|\Delta t} \rightarrow FIT_{TDDB|\Delta t|Goal} = FIT_{TDDB|Lifetime} \cdot \frac{\Delta t}{Lifetime}$$

FIT_{TOTAL} vs. VDD, TEMP

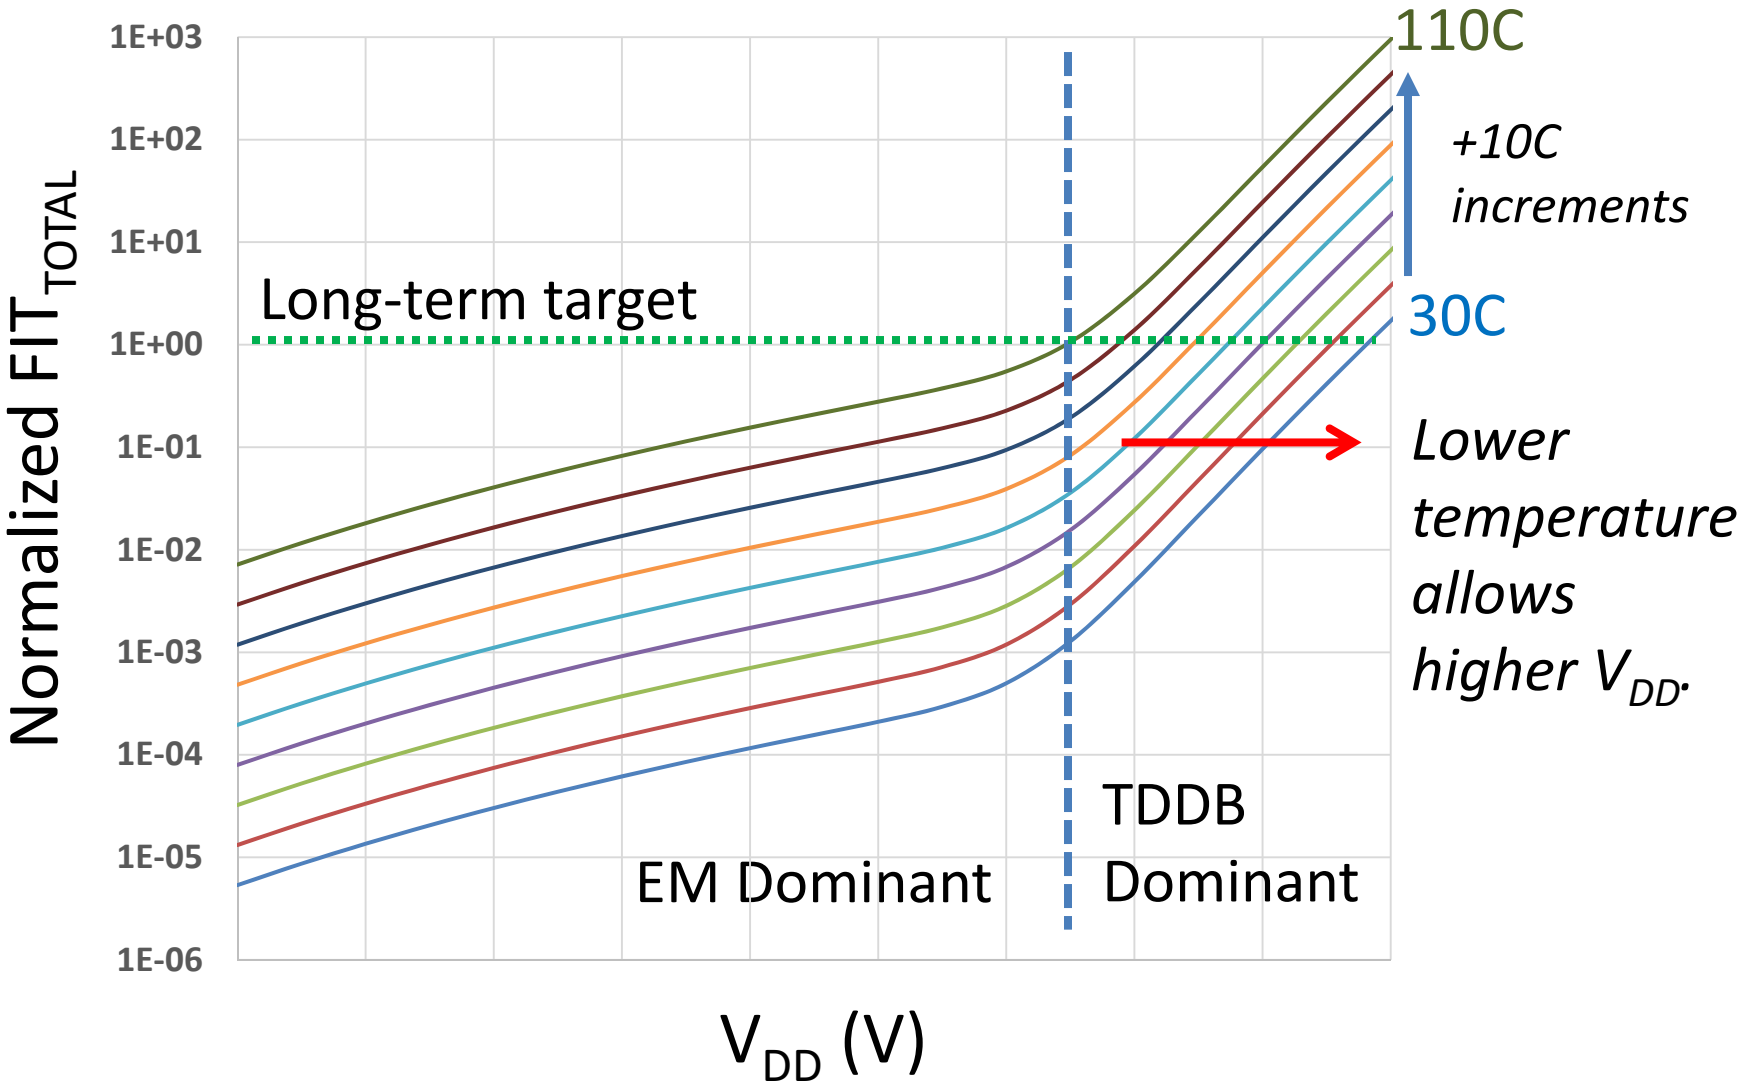

Reliability Equations

- ΔT Sample Time: 1msec
  - Smaller than P-State changes (V_{DD} constant)
  - Smaller than silicon thermal time constant (T constant)

- The incremental failure rate is curve-fit from foundry reliability models, (for ΔT =1msec), as a function of T, V_{DD}:

$$FIT_{1msec} = \underbrace{k_{TDDB} \cdot \exp^{(b_{TDDB} \cdot T)} \cdot V_{DD}^{v_{TDDB}}}_{FIT_{TDDB|1msec}} + \underbrace{k_{EM} \cdot \exp^{(b_{EM} \cdot T)} \cdot V_{DD}^{v_{EM}}}_{FIT_{EM|1msec}}$$

- FIT_{TDDB} matches foundry models extremely well.
- FIT_{EM} matches less well, mainly limits max temperature.

Filtering

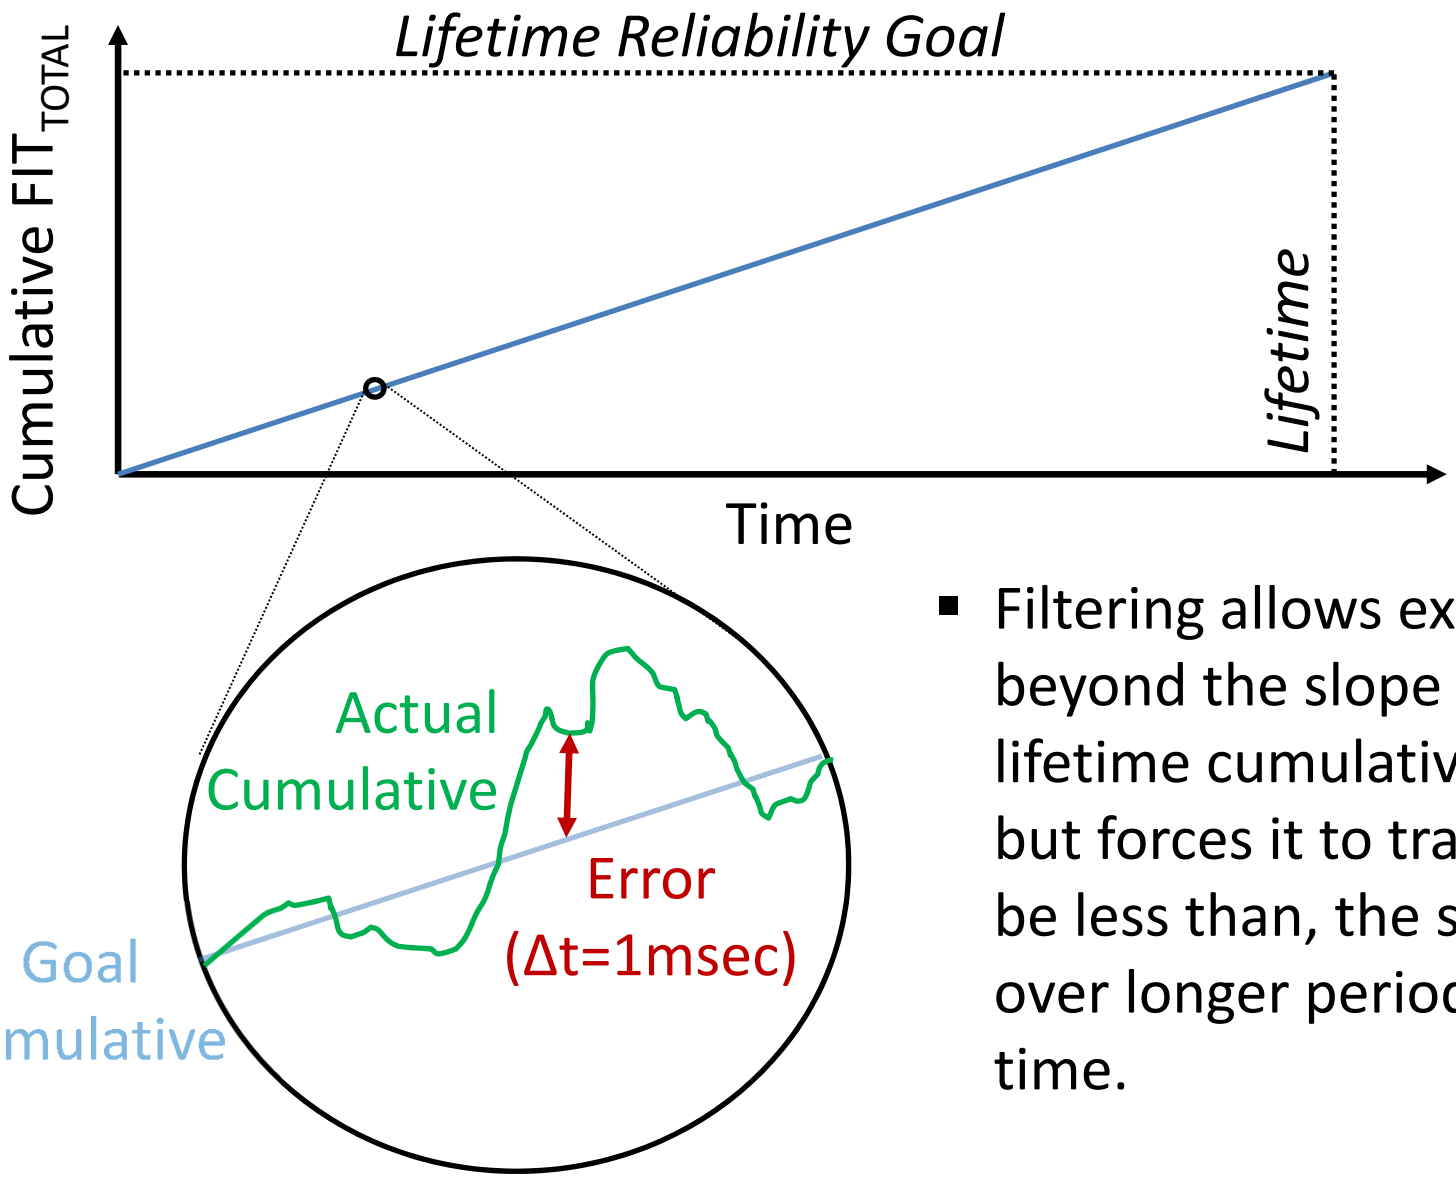

- Filtering allows excursions beyond the slope of the lifetime cumulative goal, but forces it to track to, or be less than, the slope over longer periods of time.

Filtering (cont)

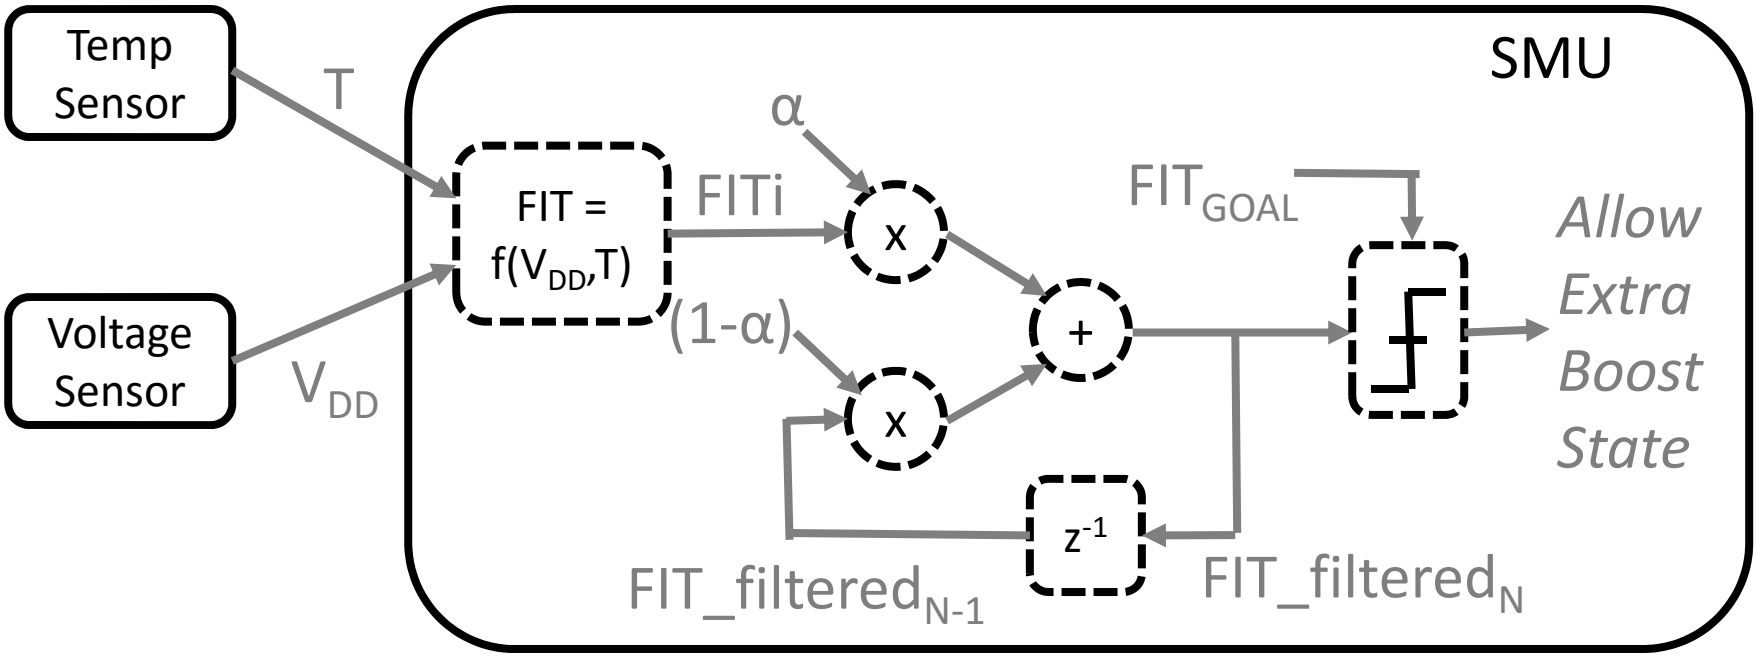

- Implemented in System Management Unit (SMU) firmware.
- α tuned to best balance peak performance vs. maximum excursion from FIT_{GOAL}.

Benefit to Typical User

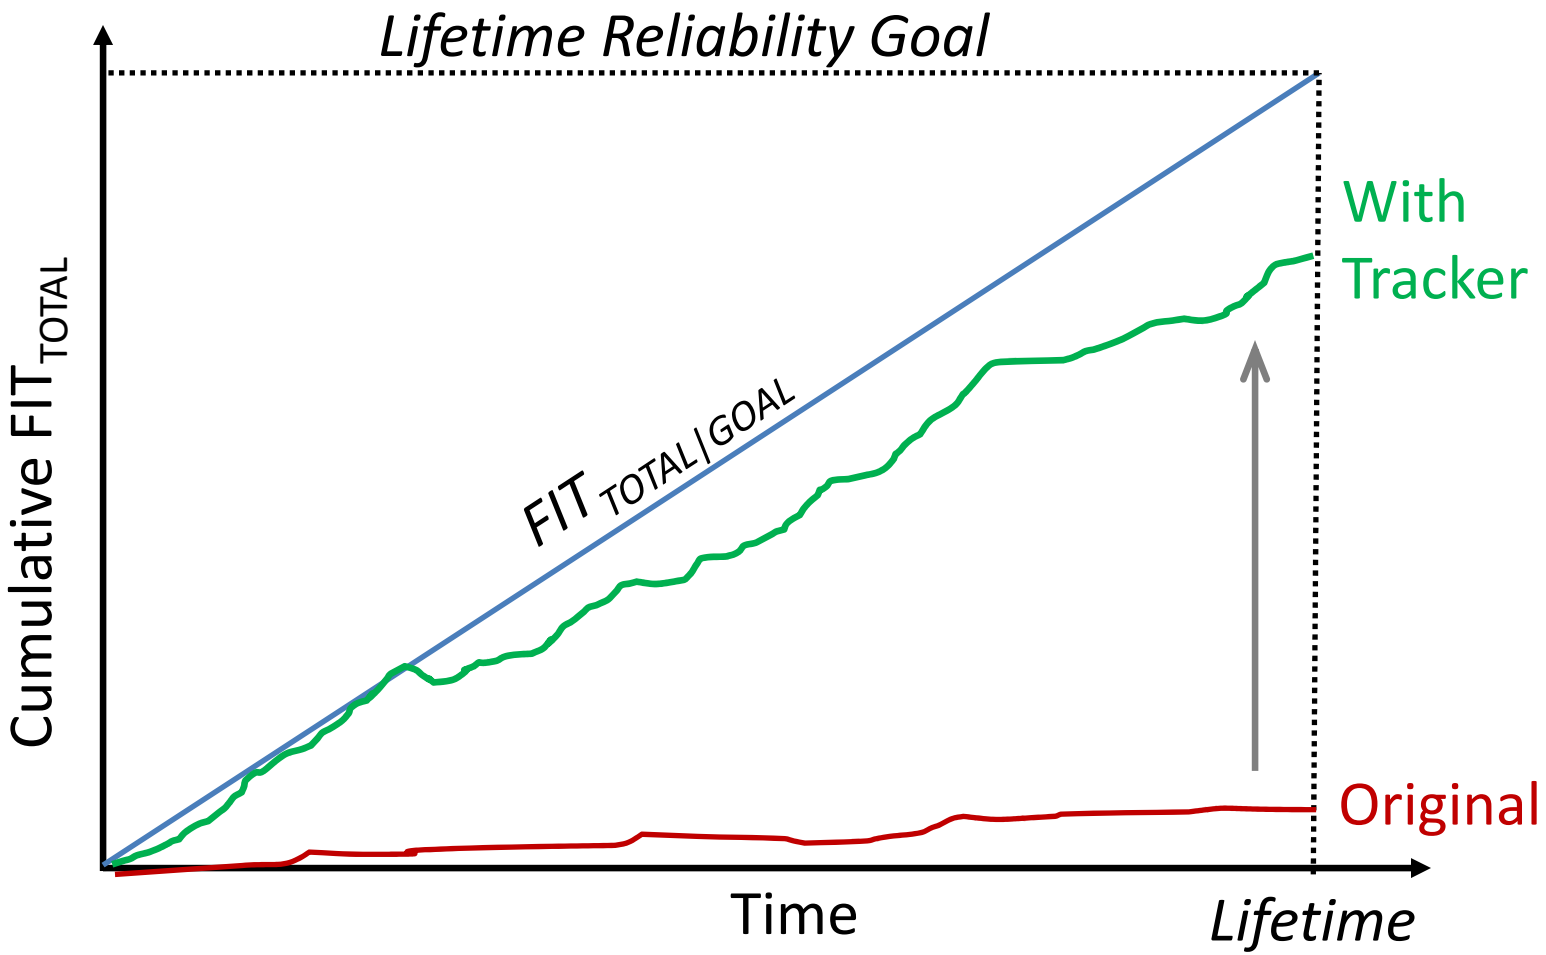

- By reducing conservative design margins, the Tracker can deliver higher boost frequencies, and deliver a better user experience.

DVR SYSTEM ARCHITECTURE

- Simple digital controller running on a 100 MHz clock
- Z-domain model of the controller implemented to optimize the system performance (settling time, ripple etc.)
- Controller logic is ~1000 sq. um
- Regular power-gating headers are re-purposed to do voltage regulation
- 48 possible wake states; 48-bit thermometer code implementation to drive the headers
- No explicit caps on VDDCORE

DVR SYSTEM ARCHITECTURE

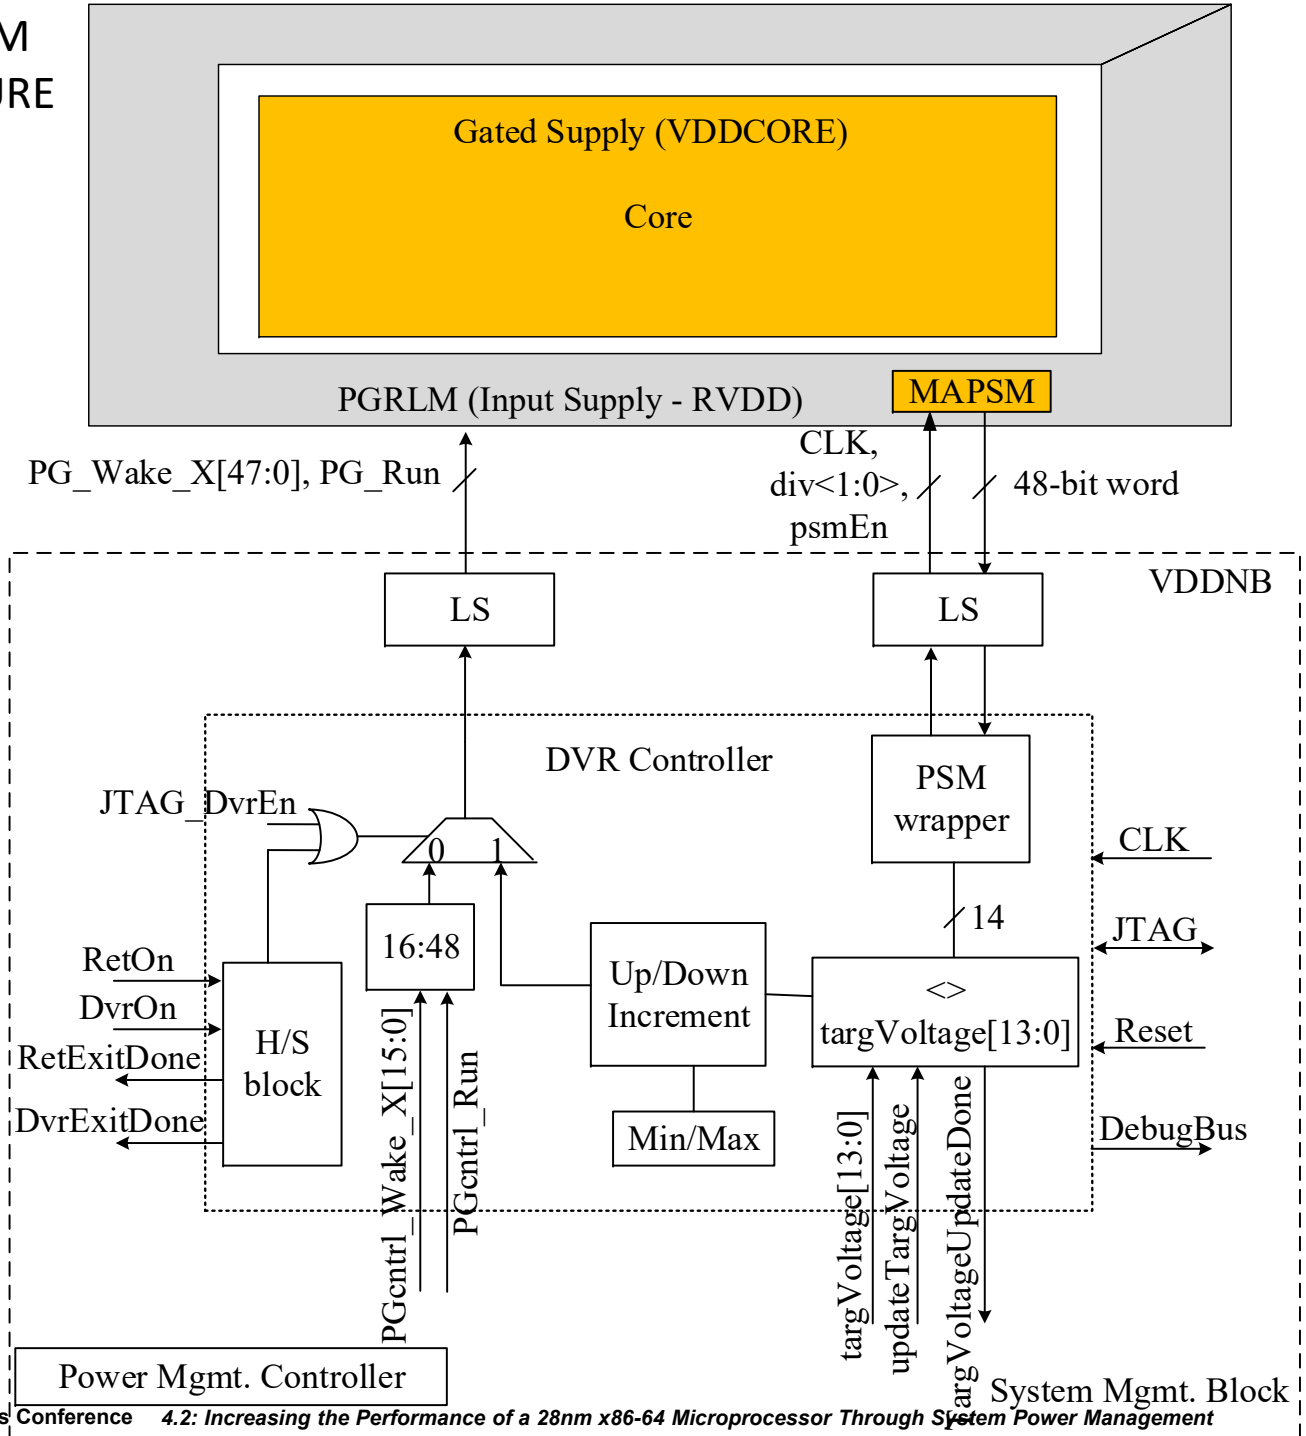

DVR CIRCUIT MODEL

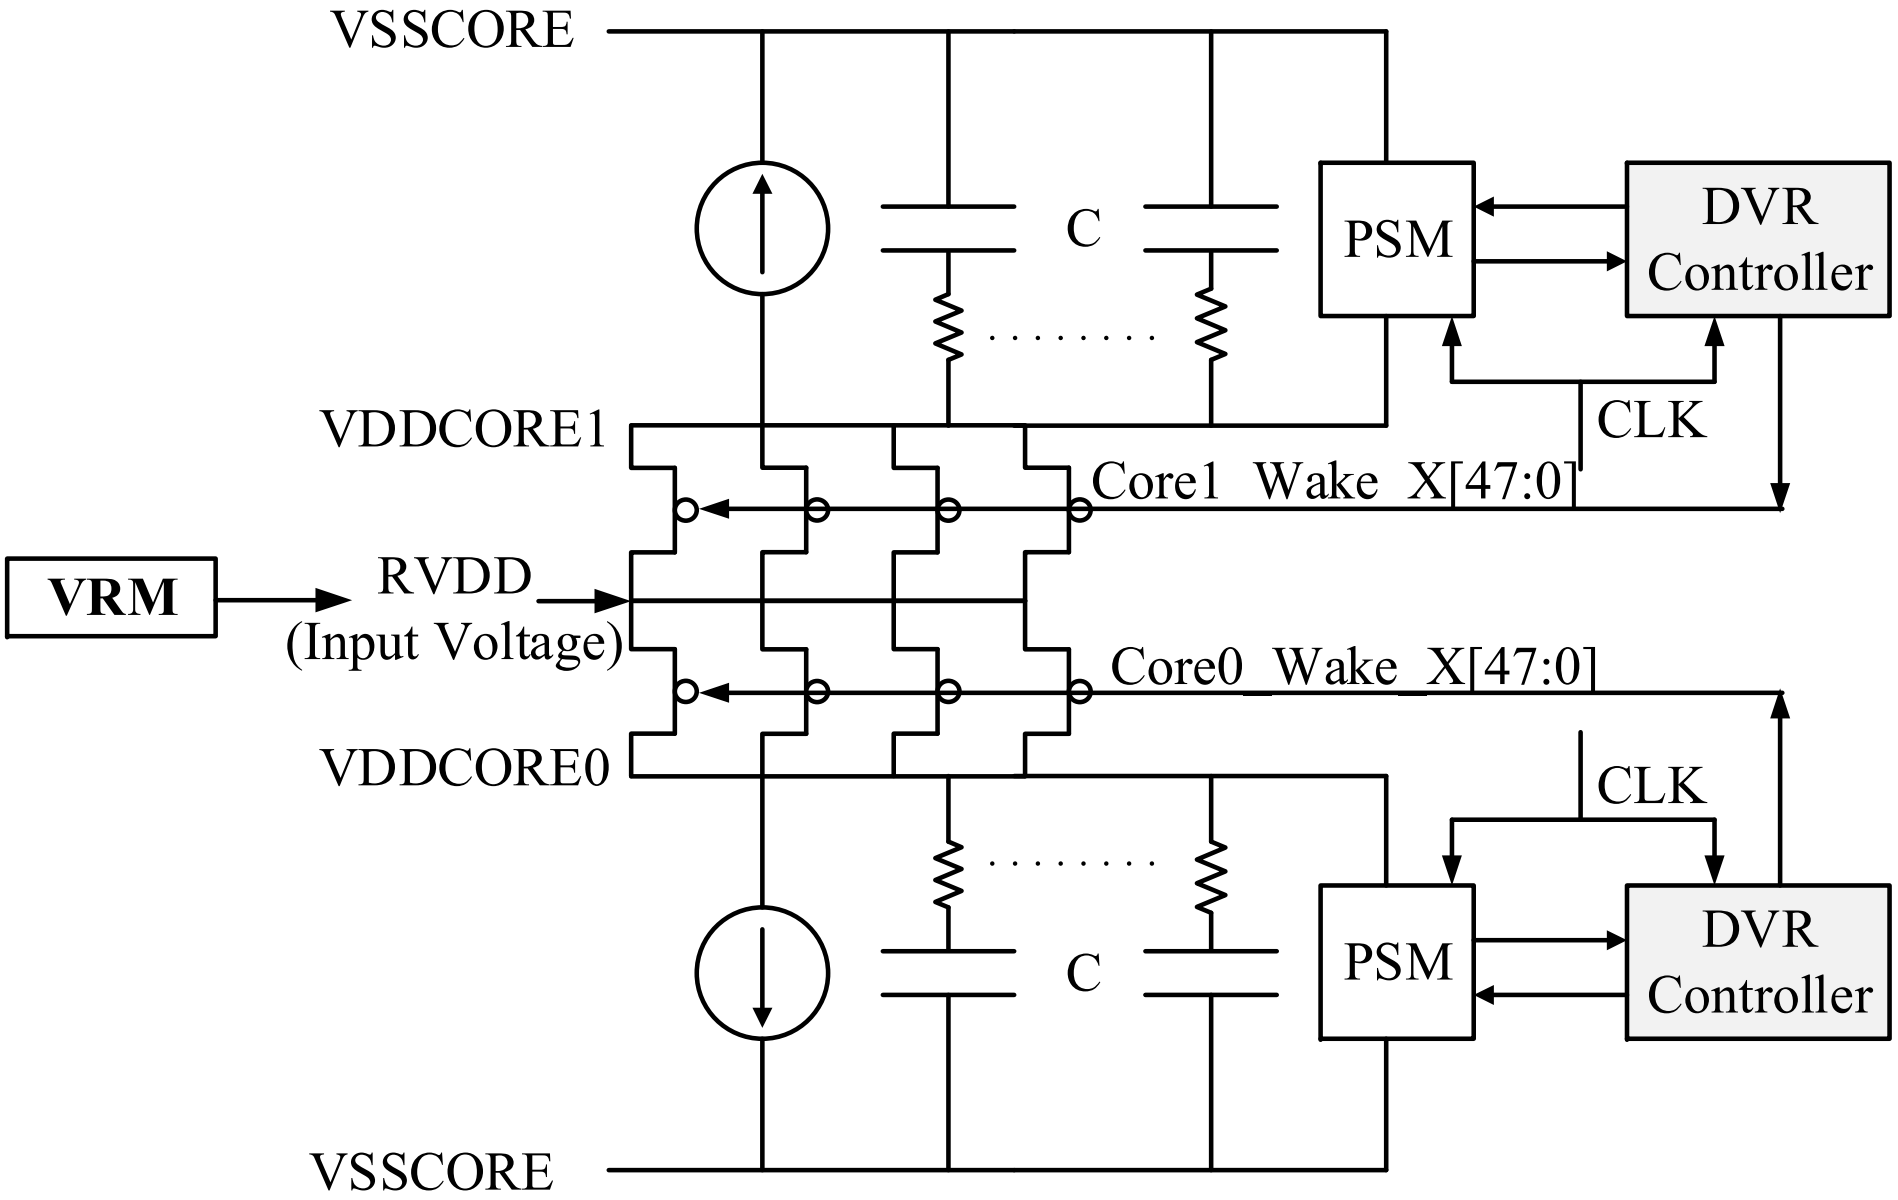

Architectural throttles

- LDO controller runs on 100MHz clocks
- The system doesn't have enough caps on regulated node to sustain large droops due to di/dt. Affects V_{min} in LDO mode.
- Employed architectural throttles to mitigate the current load, thus improve V_{min}
- V_{min} measurement by sweeping the architectural throttlers

Mode	Description
LDO-Default	No throttling
LDO-Opt1	Disable Branch Prediction
LDO-Opt2	Disable Branch Prediction + Single Issue
LDO-Opt3	Disable Branch Prediction + Single Issue + Non-Spec Execution
LDO-Opt4	Disable Branch Prediction + Single Issue + Non-Spec Execution + Disable IC + Disable DC
VRM-Default	LDO disabled, default VRM mode

Architectural throttles

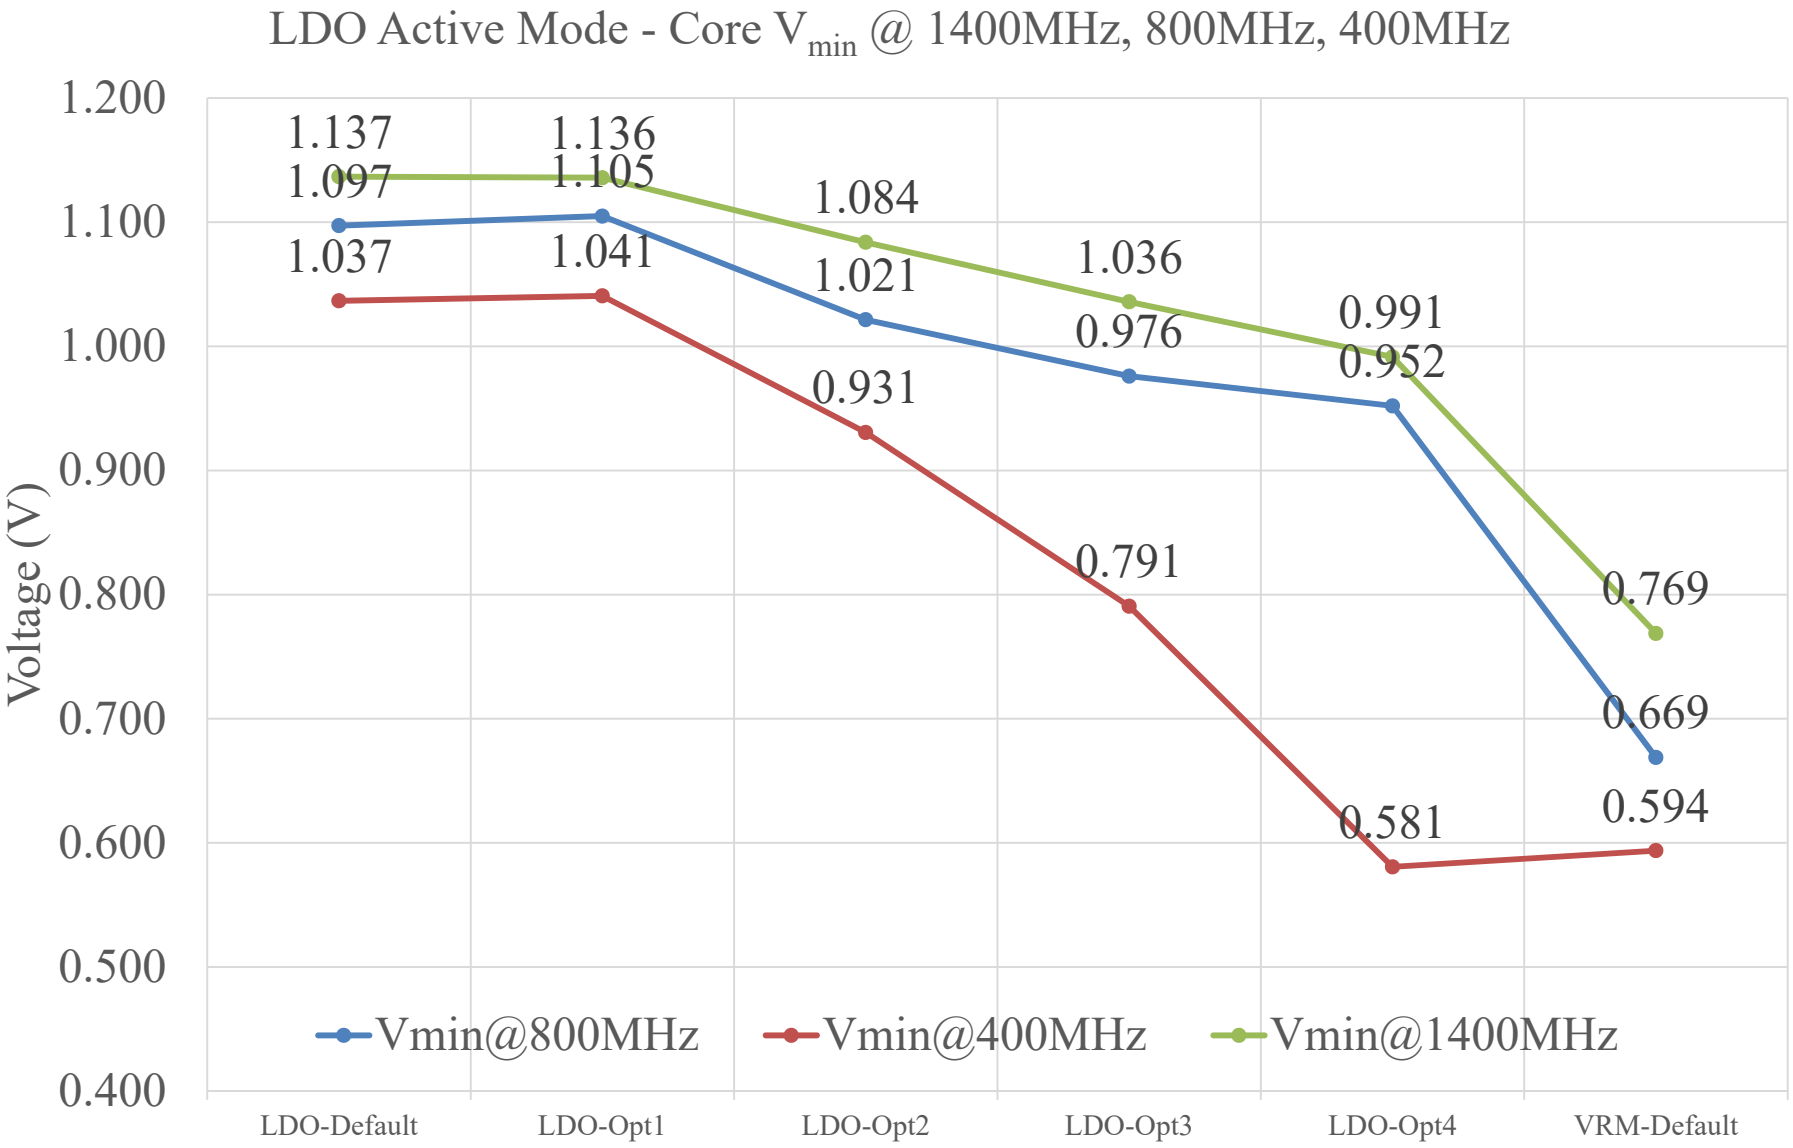

Resistance across wake states

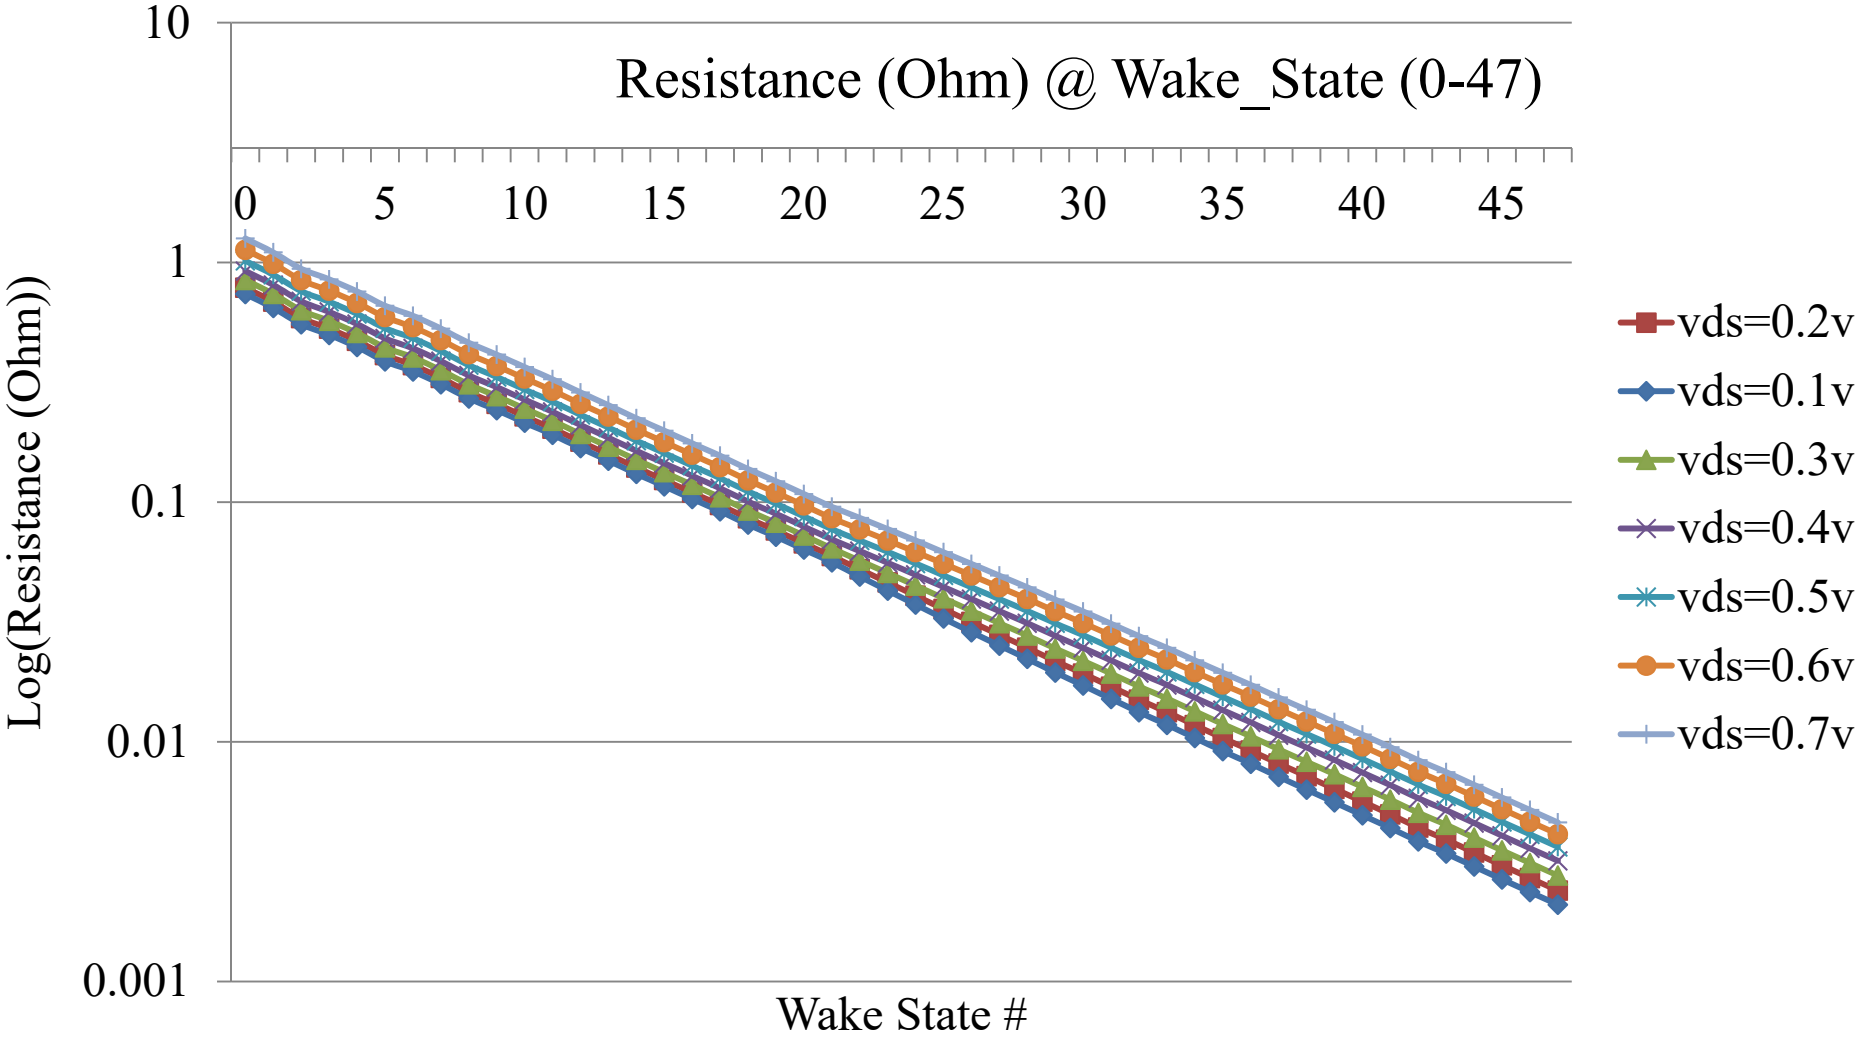

Active mode : on-die voltage measurement

- The following diagram is a measurement of on-die voltage with one core idle and another running

- The active core is running a frequency virus pattern. Amount of ripple seen is expected as the controller is running on 100MHz clock and system doesn't have enough capacitance on the regulated node.

- Based on the learning from this on-die measurement we implemented architectural throttles to mitigate the ripple and achieve low V_{min}

Retention mode : On-die voltage measurement

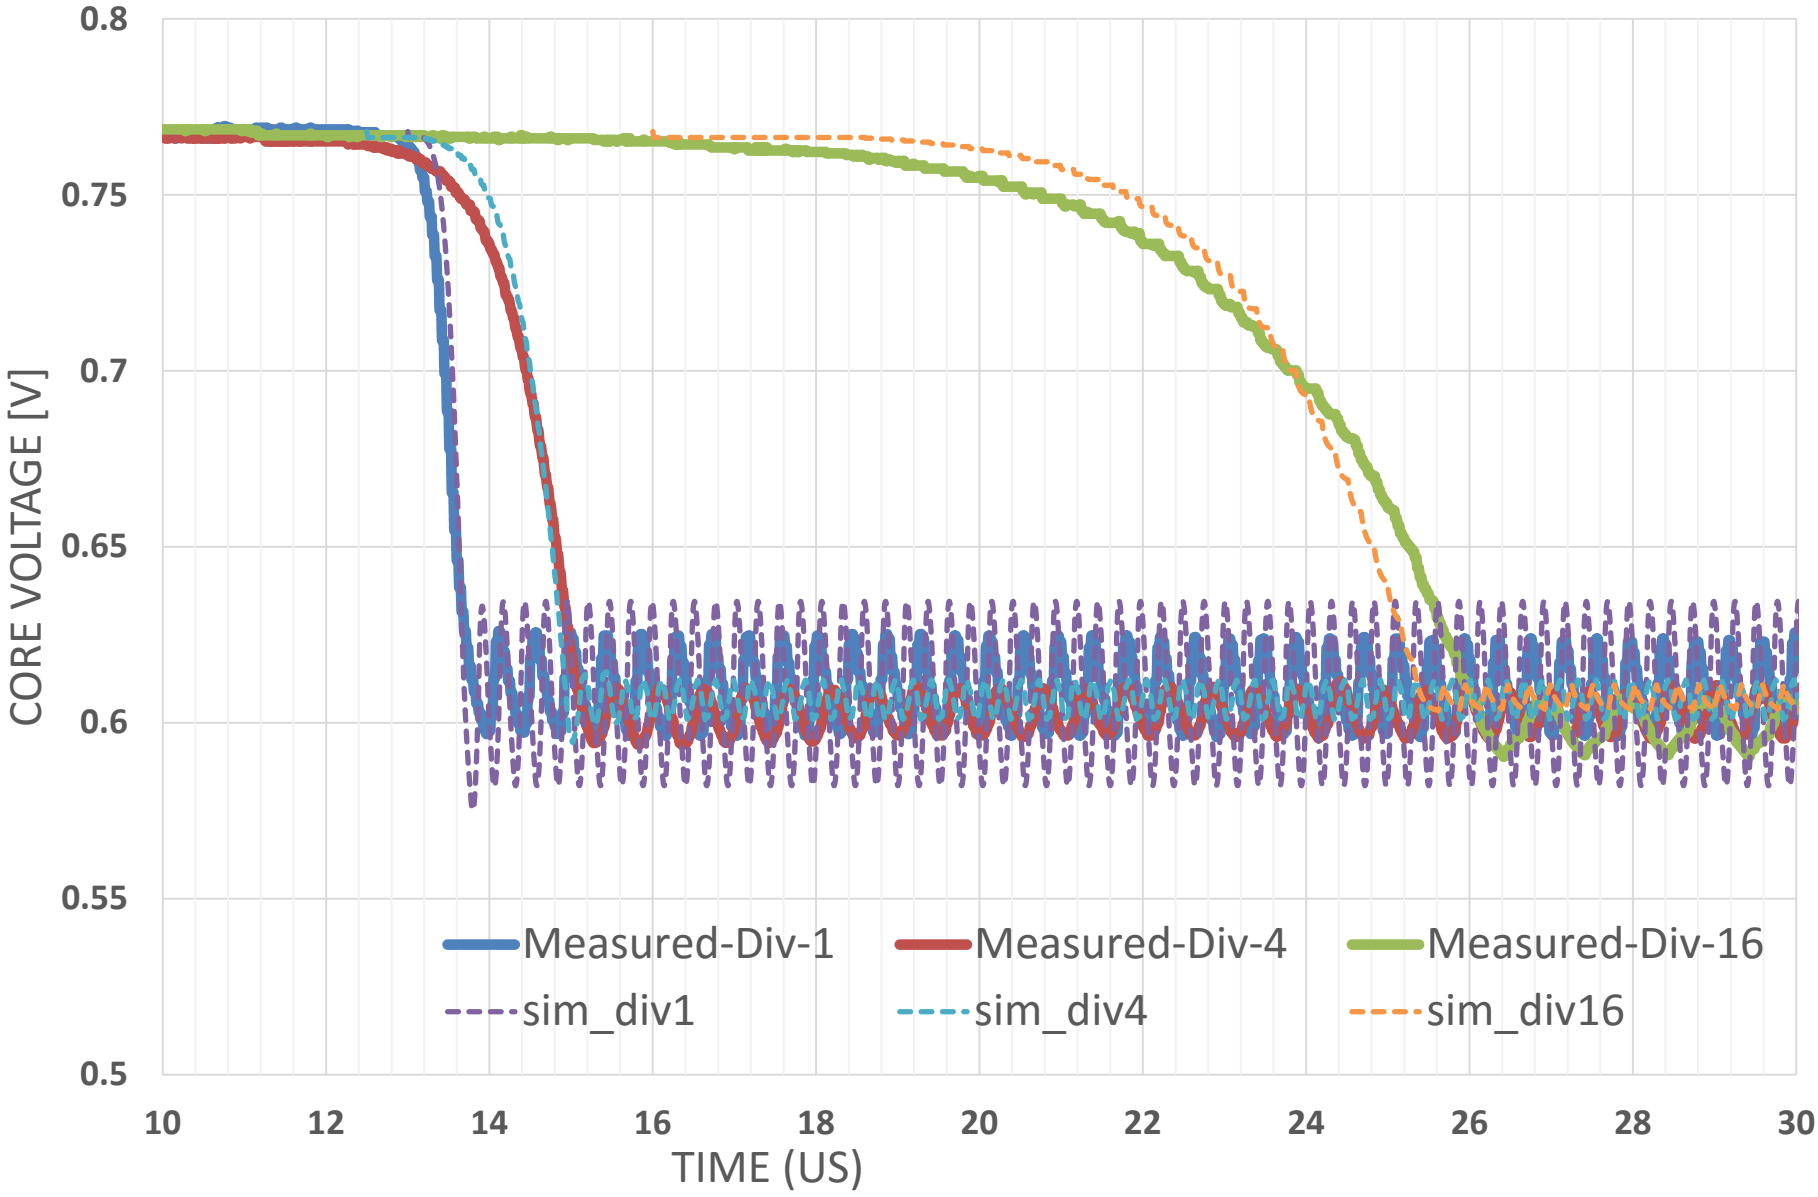

Skin Temperature Aware Power Management

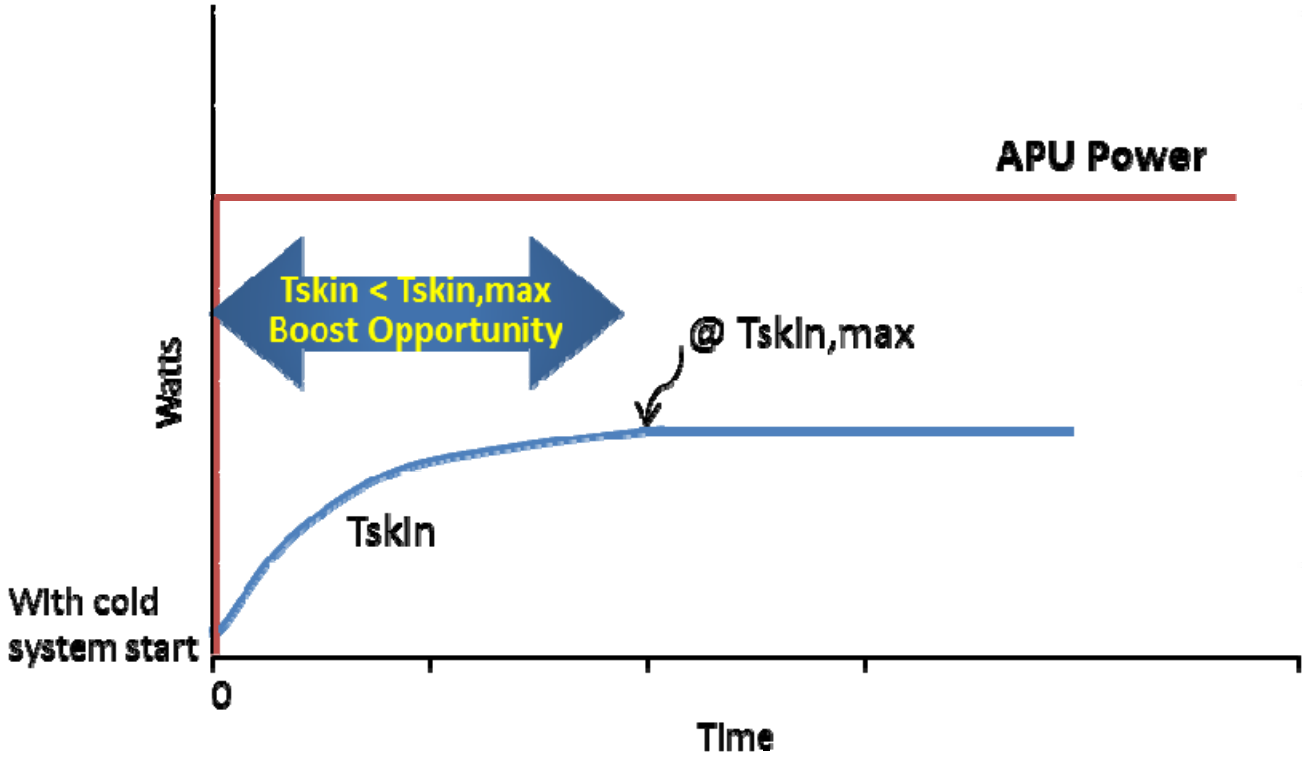

Thin form-factor notebook power is typically limited by steady state skin temperature. This steady state power level, when consumed indefinitely, will cause the chassis skin temperature to reach a specified skin temperature limit.

Due to the thermal capacitance of the system it can take several minutes to reach the skin temperature limit from a "cold" idle state.

Skin Temperature Aware Power Management

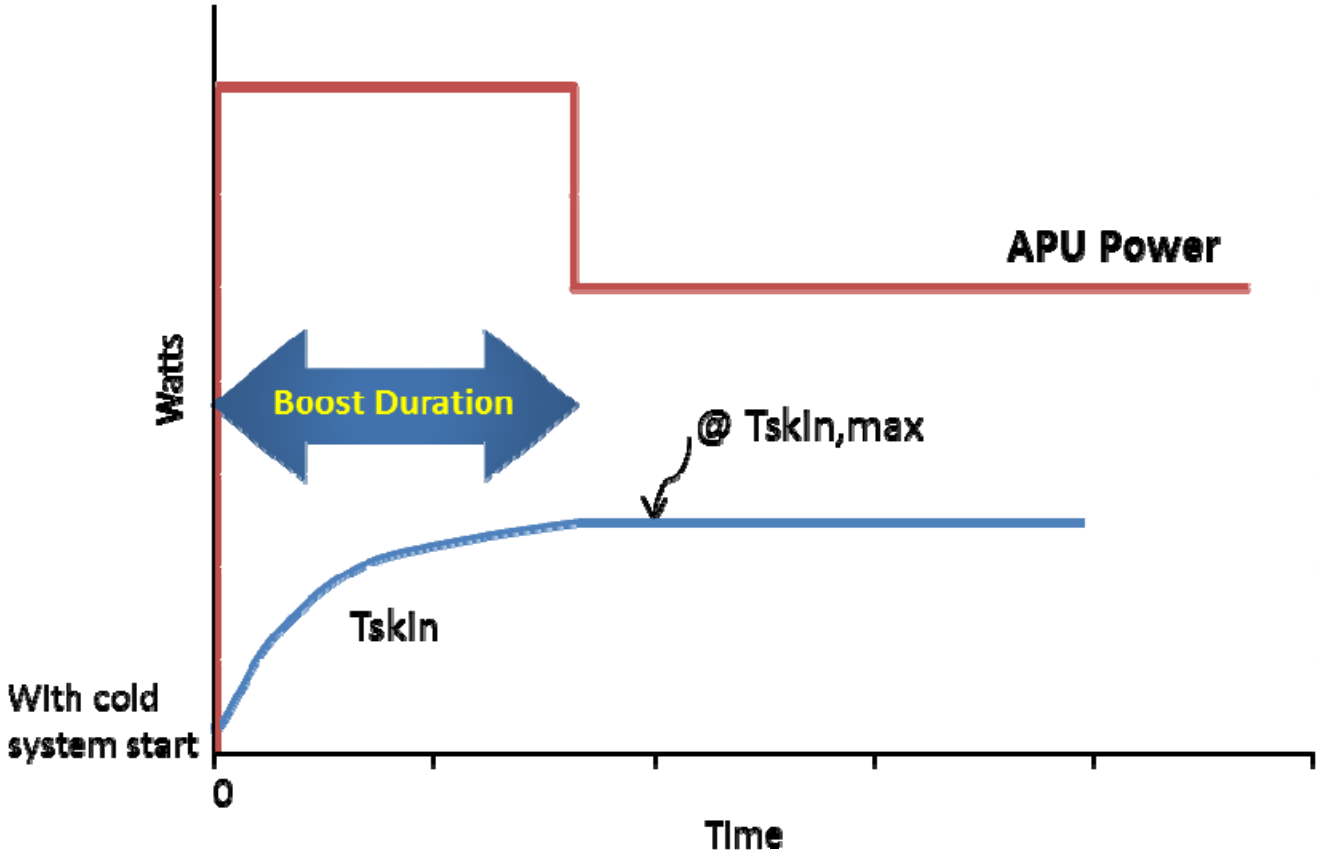

STAPM (Skin Temperature Aware Power Management) improves performance by boosting APU CPU and/or GPU frequencies for as long as the estimated platform skin temperature remains below the specified limit.

STAPM Controller

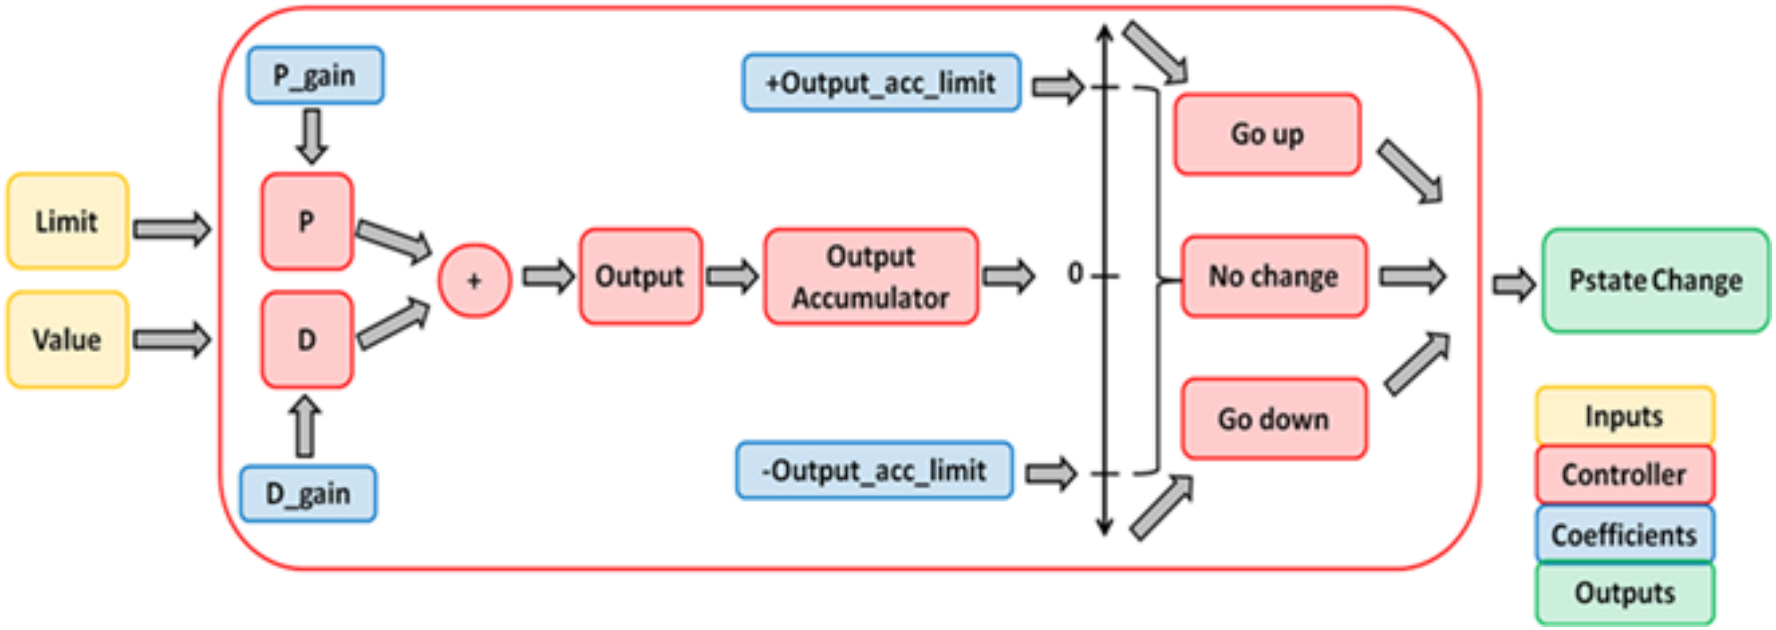

- BR STAPM models the thermal capacitance of the platform using a simple alpha filter applied to the calculated APU power.
- This alpha filtered APU power is proportional to skin temperature.
- Alpha filter coefficients can be optimized by characterizing the chassis skin temperature response to APU power.

STAPM Controller

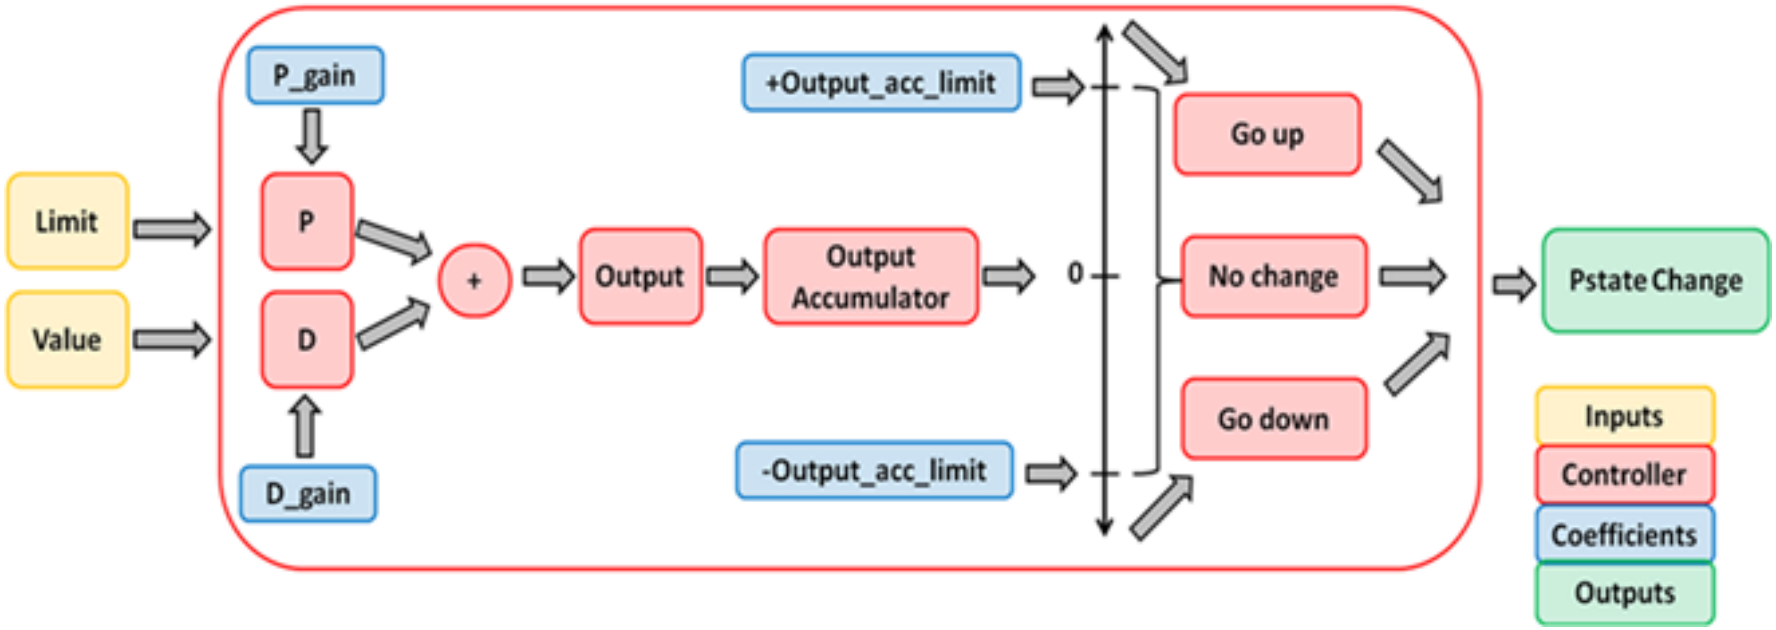

- In mission-mode, skin temperature is controlled by controlling APU power with a Proportional + Derivative + Accumulator algorithm in system management firmware.
- Control loop coefficients are selected to optimize boost duration, limit skin temperature overshoot, and optimize the user experience by gradually reducing boost as the loop controls to the skin temperature limit.

STAPM: 3DMark11 Power and Frequency Profiles

APU Power (CPU + GPU), Green
Alpha Filtered APU Power, Black
Tskin Power Limit, dashed Black
GPU Power, Pink
CPU Power, Blue

CPU Frequency, Blue
GPU Frequency, Pink

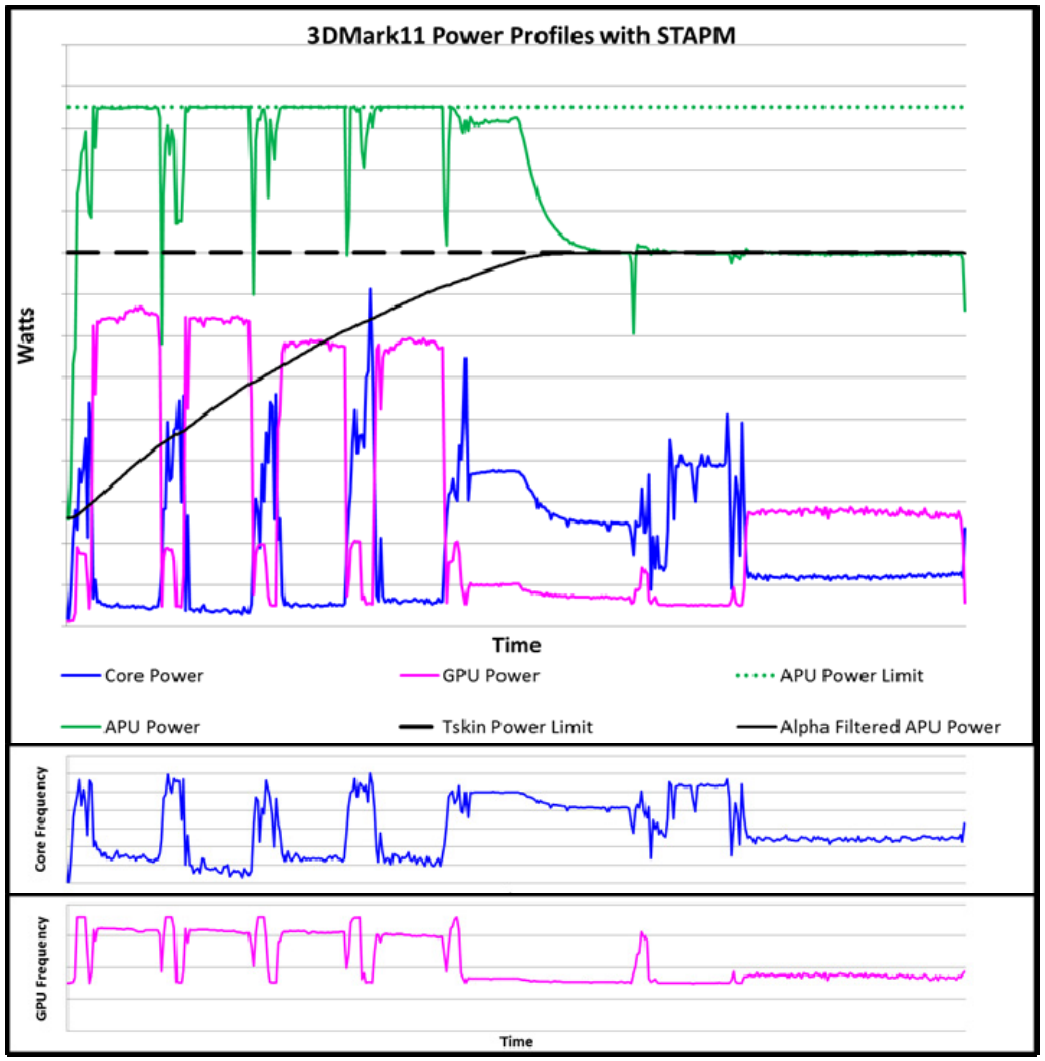

Benchmark results with STAPM Enabled

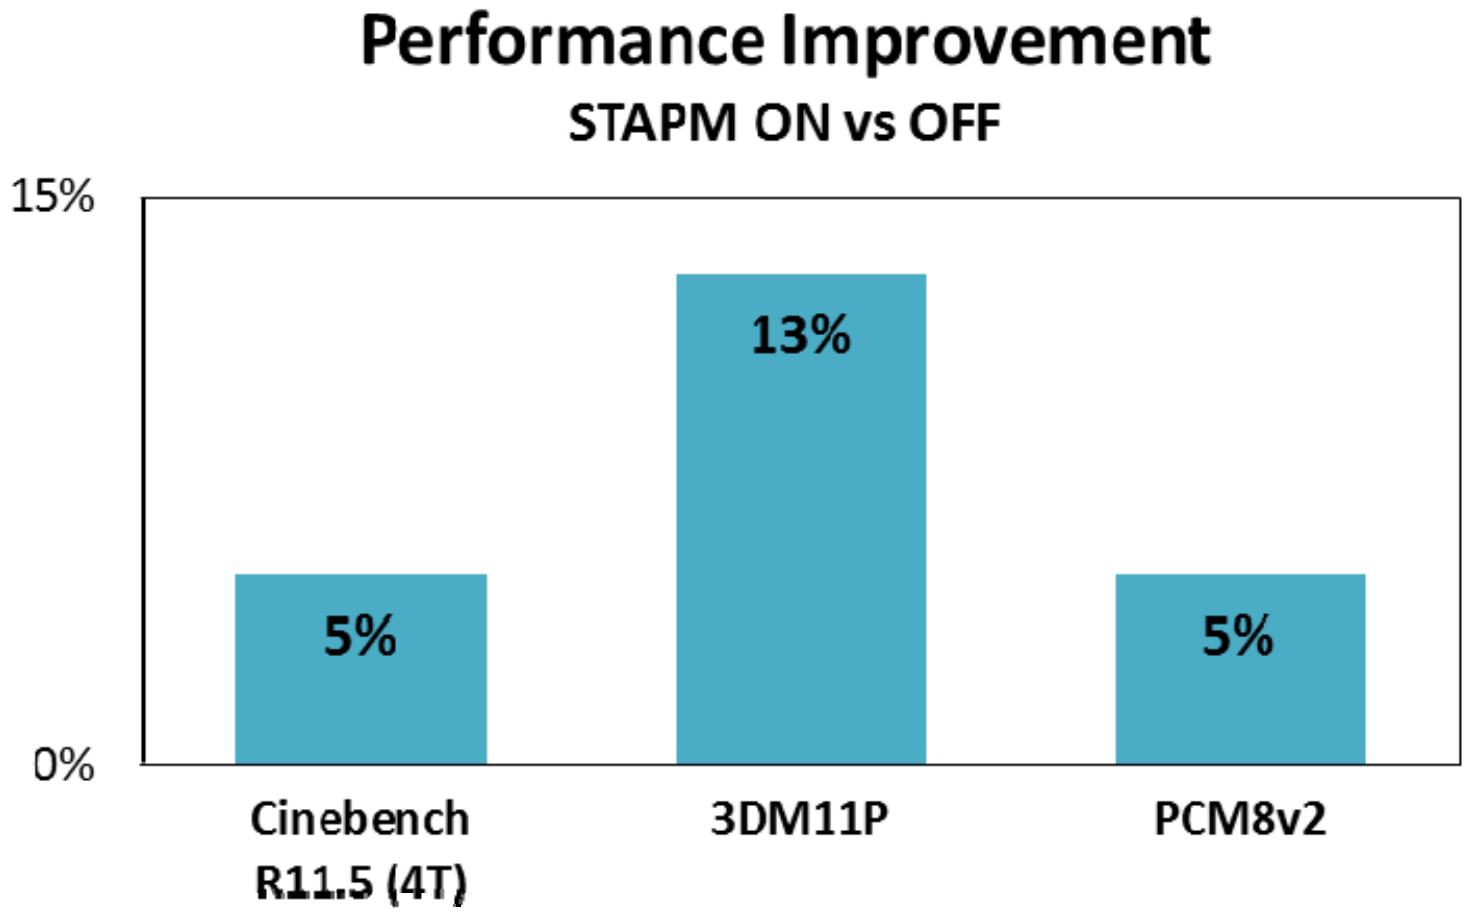

Boot Time Power Supply Calibration

- In a traditional or even variable voltage binning scheme, an offset is applied between ATE speed binning and final system level testing. This compensates for the power delivery differences between the 2 environments.

- Additionally, gaurdbands are added for:
  - Voltage regulator tolerance
  - Aging to compensate for transistors slowing down over time

- BTC compensates for all 3 of these items allowing for a safe, but more aggressive voltage for each p-state resulting in reduced power.

DC BTC – Calibrate at Temp Inversion Voltage

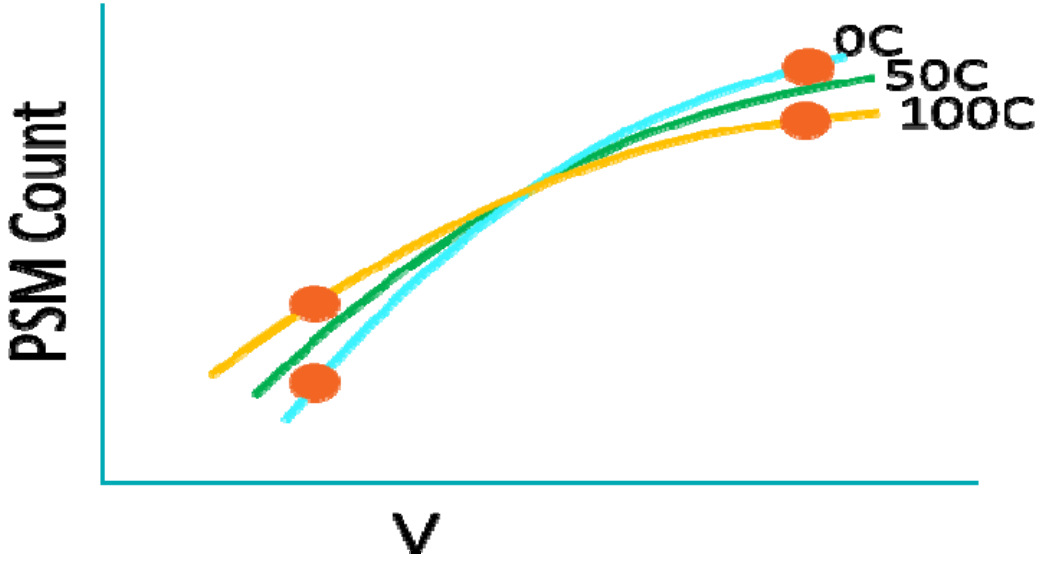

- At boot up, we temporarily place the IP voltage at the inversion voltage when reading the values from the PSMs to remove any temperature error.
- The PSM values are compared against the values obtained on ATE and the difference is converted to a voltage guard band based on pre-characterized (on ATE) PSM vs Voltage curves.
- This removes the DC tolerance regulator guardband as the DC accuracy of the voltage supply has been directly measured.

BTC Part Specific ATE to System Offset

- BTC calculates a per-part Tester to real system offset.
- As part of Si characterization, the minimum PSM (power supply monitor) count is captured at the ATE insertion as well as in a real system.
- This data is used to derive a part-specific “AC” voltage offset between the 2 platforms based on each part’s Sidd, Cac, operating tempature, individual p-state (frequency), and measured AC droop.
- Tester2System_Voltage_Offset = (Sidd * Coeff_Sidd) + (Cac * Coeff_Cac) + (Temperature * Coeff_Temp) + (Frequency * Coeff_Frequency) + (AteVdroop * Coeff_ateVdroop) + sigma_for_yield

BTC Used to Reduce Aging Margin

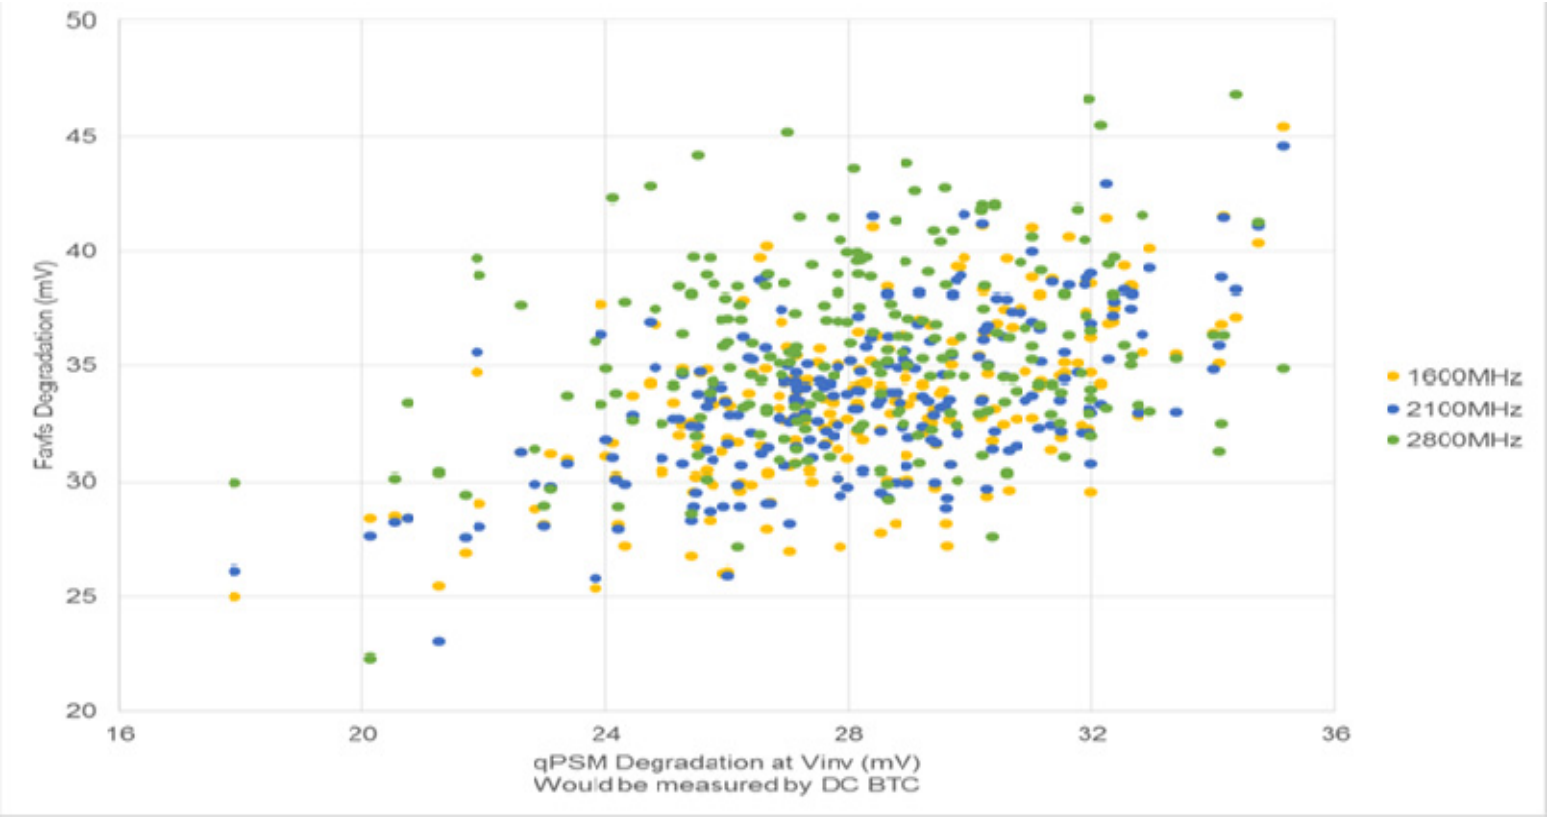

- Since BTC aligns the PSMs to generate the voltage, the BTC algorithm will naturally compensate for aging effects.
- As the part slows down over the time, the PSMs will report that they need more voltage at boot time
- As shown in the chart above, the PSM count tracking covers about 2/3 of the needed aging guardband allowing that voltage to be removed to save power.

Conclusion

- Silicon and platform performance is ultimately limited by thermal, voltage, and current constraints.
- Additional reliability guardbands further constrain performance
- Power management techniques can be applied to reduce the impact of these performance limitations.
- When all the techniques were added together, the Bristol Ridge APU had an increase of 10-15% on most benchmarks.

A 20nm 2.5GHz Ultra-Low-Power Tri-Cluster CPU Subsystem with Adaptive Power Allocation for Optimal Mobile SoC Performance

Hugh T Mair, Gordon Gammie, Alice Wang, Rolf Lagerquist, C.J. Chung, Sumanth Gururajarao, Ping Kao, Anand Rajagopalan, Anirban Saha, Amit Jain, Ericbill Wang, Shichin Ouyang, Huajun Wen, Achuta Thippana, HsinChen Chen, Syed Rahman, Minh Chau, Anshul Varma, Brian Flachs, Mark Peng, Alfred Tsai, Vincent Lin, Ue Fu, Wuan Kuo, Lee-Kee Yong, Clavin Peng, Leo Shieh, Jengding Wu, Uming Ko

MEDIATEK

Outline

- SoC Overview
- Tri-Cluster CPU Subsystem
  - LP Cluster Power-efficiency/Performance benefits
- Adaptive-Power-Allocation (APA)
  - APA Modes: APA-CD and APA-MP
  - Power Meters: Dynamic and Leakage
  - Silicon results
- SupplEyeScan (SES)
  - Motivation
  - Design Overview and Operating Modes
  - Silicon results
- Summary

SoC Overview

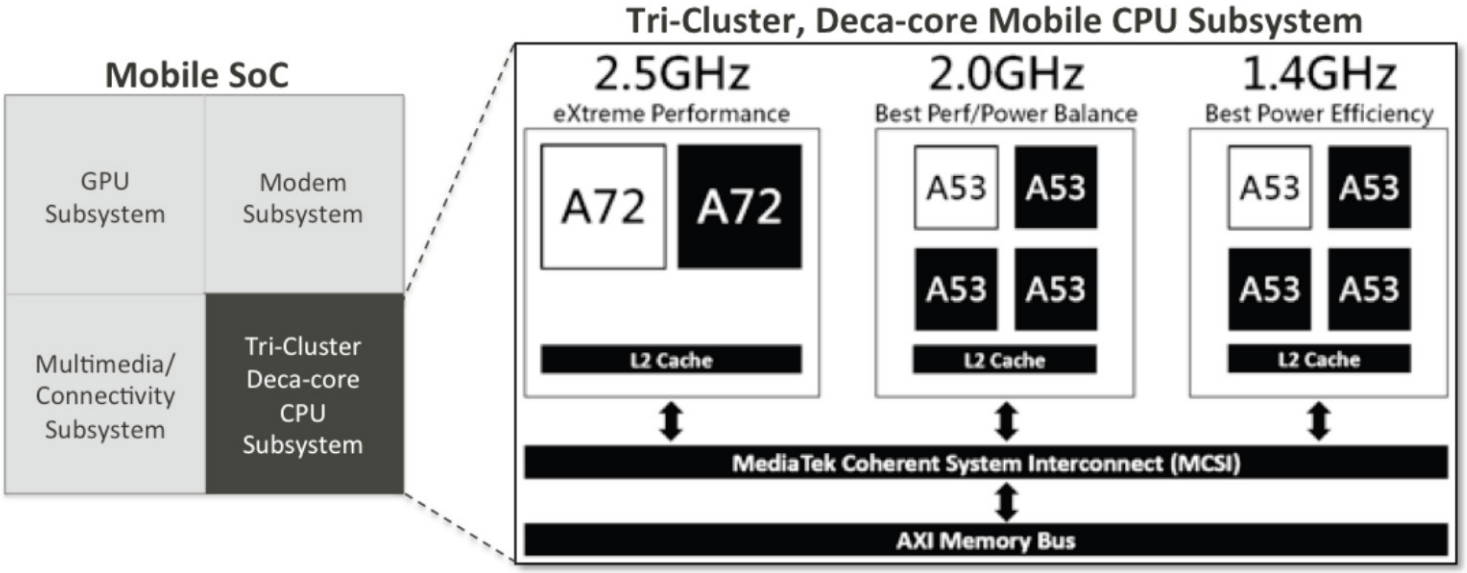

- **Heterogeneous ARM-v8A Tri-Cluster Deca-Core CPU**
  - **2.5GHz 2x ARM CortexA72**
  - **2.0GHz 4x ARM CortexA53**
  - **1.4GHz 4x ARM CortexA53**
- **GPU Subsystem**
- **Multimedia (32MPixel/24fps camera support)**
- **802.11ac, GPS**
- **Modem: LTE FDD/TDD R11 Cat-6 with 20+20 carrier aggregation (300/50Mb/s) DC-HSPA+, TD-SCDMA,Edge, CDMA2000 1x/EVDO Rev. A (SRLTE)**

Outline

- SoC Overview
- **Tri-Cluster CPU Subsystem**
  - **LP Cluster Power-efficiency/Performance benefits**
- Adaptive-Power-Allocation (APA)
  - APA Modes: APA-CD and APA-MP
  - Power Meters: Dynamic and Leakage
  - Silicon results
- SupplEyeScan (SES)
  - Motivation
  - Design Overview and Operating Modes
  - Silicon results
- Summary

Tri-Cluster CPU Subsystem

- **ULP Cluster**
  - 4x Cortex-A53 cores
  - Max freq: 1.4GHz
  - Optimized for Ultra-Low power applications
- **LP Cluster** NEW
  - 4x Cortex-A53 cores
  - Max freq: 2.0GHz
  - +40% power-eff. vs. HP Cluster
  - +40% performance vs. ULP Cluster
- **HP Cluster**
  - 2x Cortex-A72 cores
  - Max freq: 2.5GHz
  - Optimized for high-performance

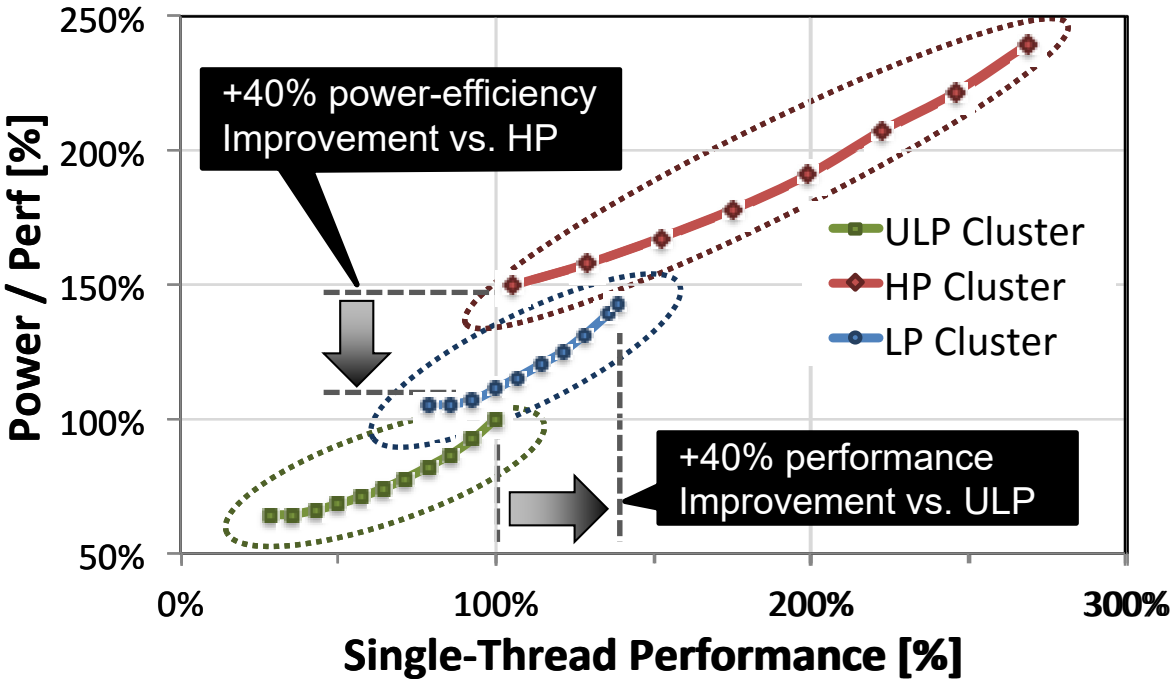

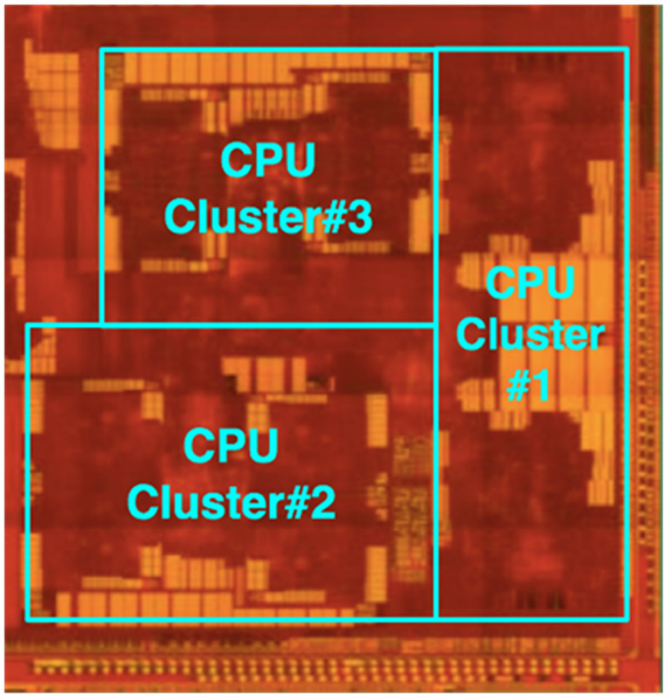

Tri-Cluster CPU Subsystem

- **ULP Cluster**
  - Light tasks

- **LP Cluster**
  - Medium tasks

- **HP Cluster**
  - Heavy tasks

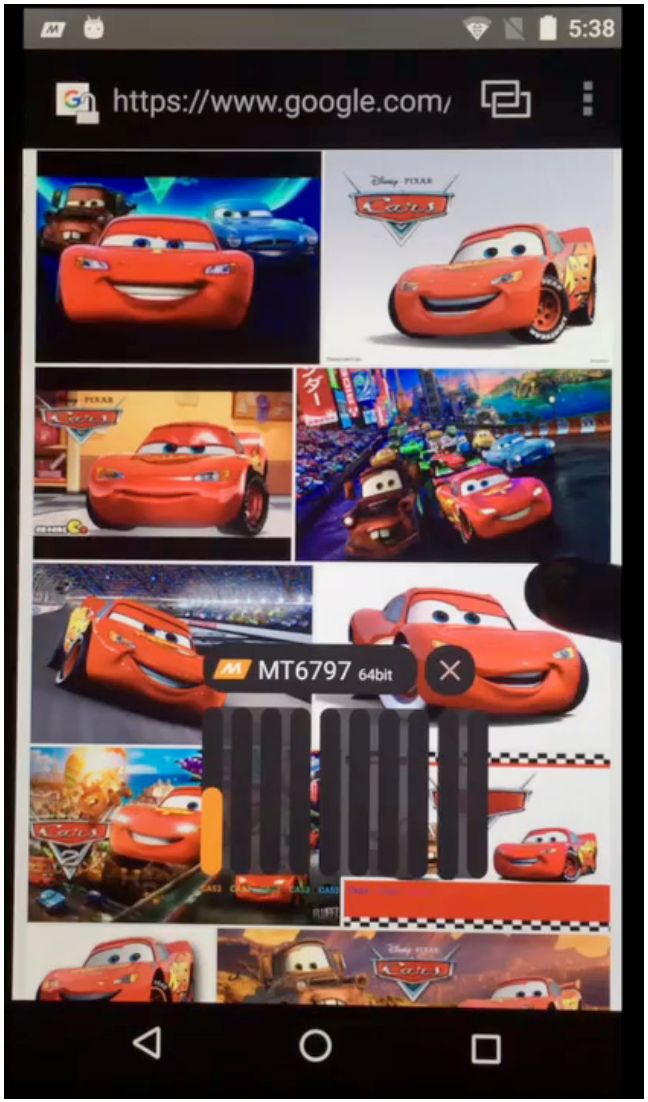

Tri-Cluster CPU Subsystem

- **ULP Cluster**
  - Light tasks

- **LP Cluster**
  - Medium tasks

- **HP Cluster**
  - Heavy tasks

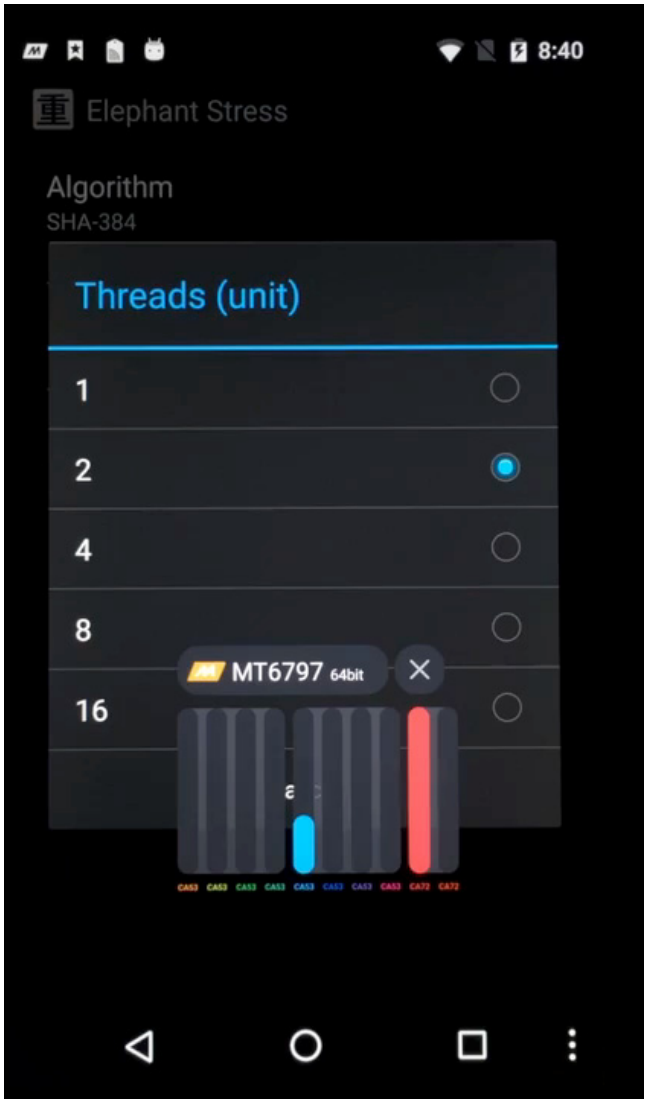

Tri-Cluster CPU Subsystem

- **ULP Cluster**
  - Light tasks

- **LP Cluster**
  - Medium tasks

- **HP Cluster**
  - Heavy tasks

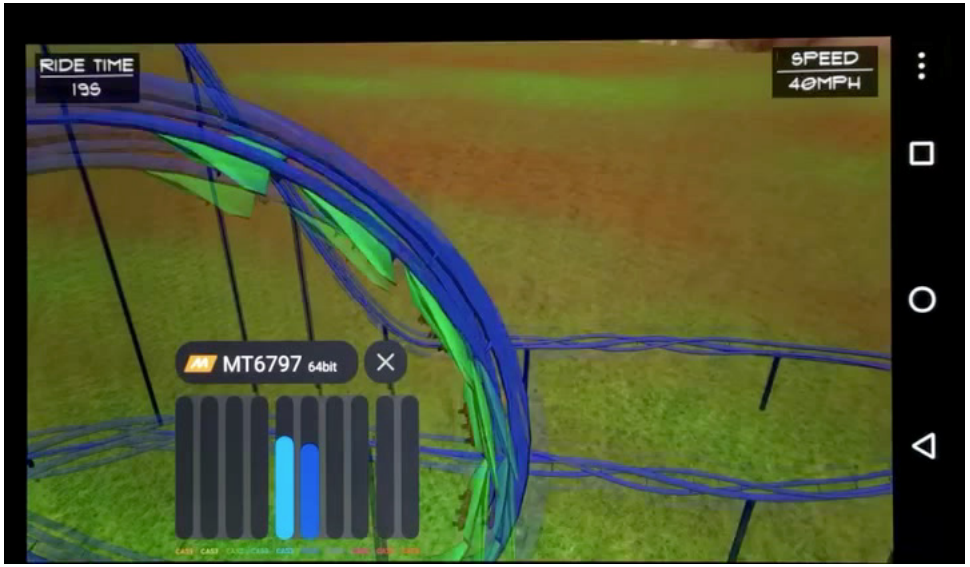

Outline

- SoC Overview
- Tri-Cluster CPU Subsystem
  - LP Cluster Power-efficiency/Performance benefits
- **Adaptive-Power-Allocation (APA)**
  - **APA Modes: APA-CD and APA-MP**
  - **Power Meters: Dynamic and Leakage**
  - **Silicon results**
- SupplEyeScan (SES)
  - Motivation
  - Design Overview and Operating Modes
  - Silicon results
- Summary

Adaptive Power Allocation (APA)

- Power supply & thermal design limit max. power
  - Clk-Dithering to hold limit
    - Reduced pwr. eff.
  - Hard Limit

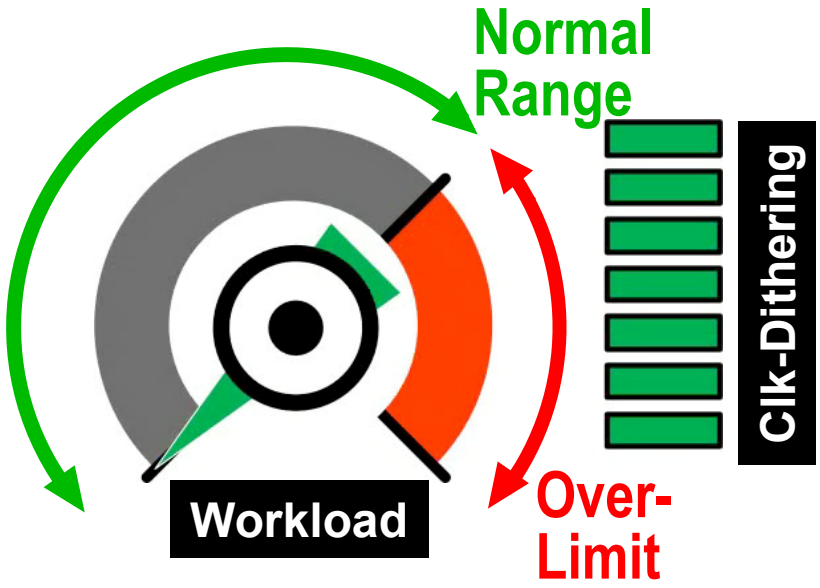

Adaptive Power Allocation (APA)

- Power supply & thermal design limit max. power
  - Clk-Dithering to hold limit
    - Reduced pwr. eff.
  - Hard Limit

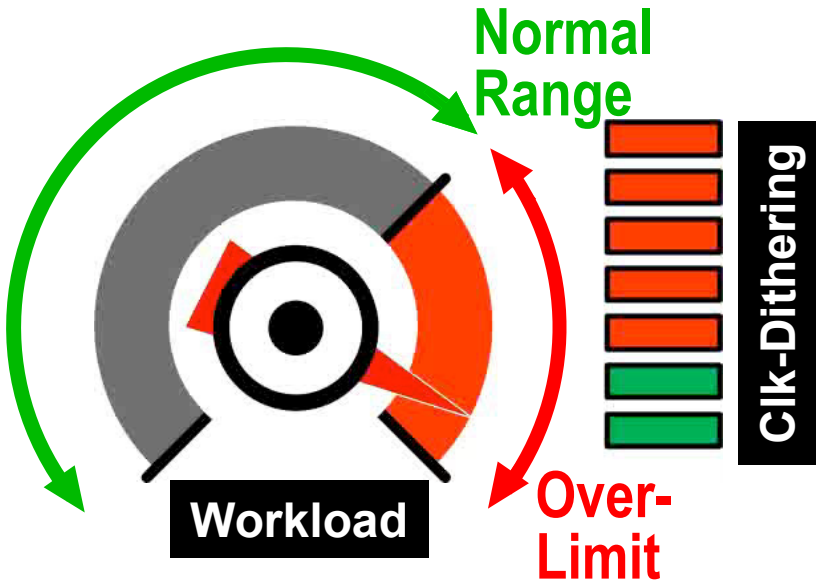

Adaptive Power Allocation (APA)

- Power supply & thermal design limit max. power
  - Clk-Dithering to hold limit
    - Reduced pwr. eff.
  - Hard Limit

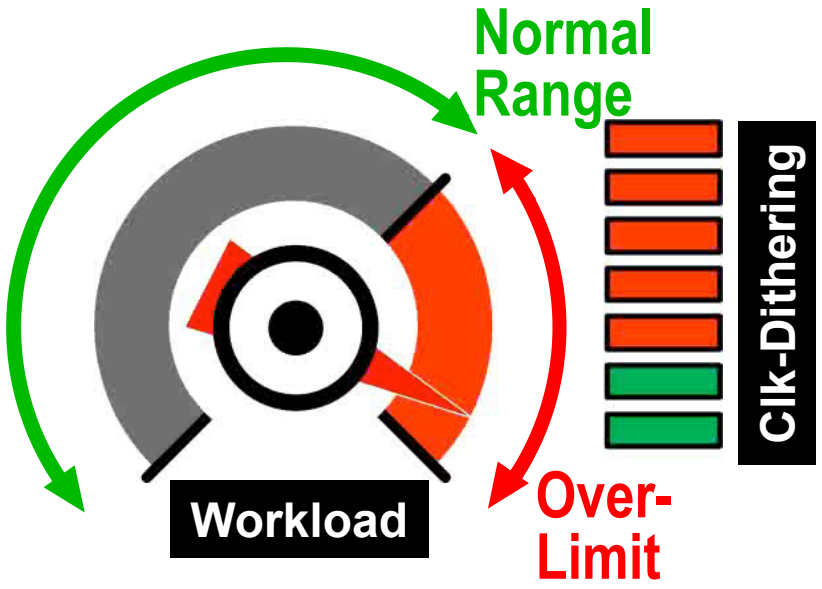

- Sharing budget across CPUs reduces Clk-Dithering
- Real-time metering:
  - Dynamic Pwr.
  - Leakage Pwr.

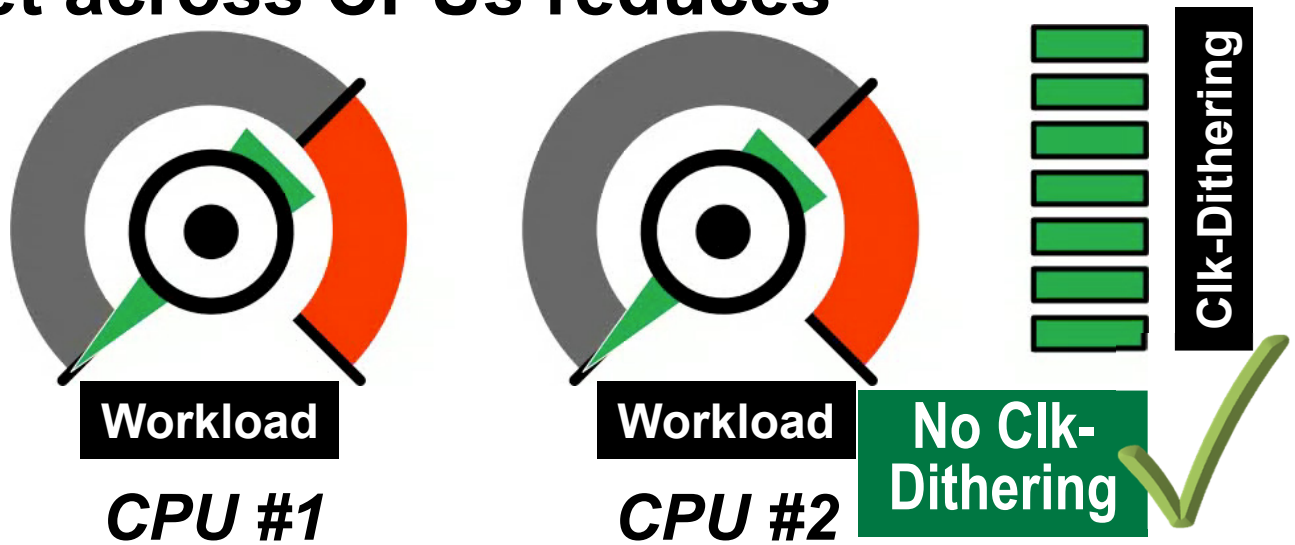

Adaptive Power Allocation (APA)

- Power supply & thermal design limit max. power
  - Clk-Dithering to hold limit
    - Reduced pwr. eff.
  - Hard Limit

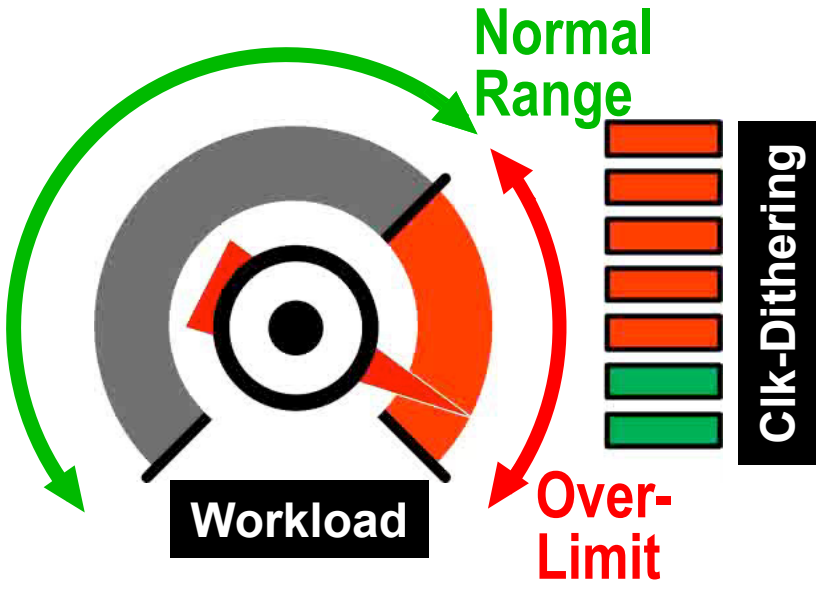

- Sharing budget across CPUs reduces Clk-Dithering
- Real-time metering:
  - Dynamic Pwr.
  - Leakage Pwr.

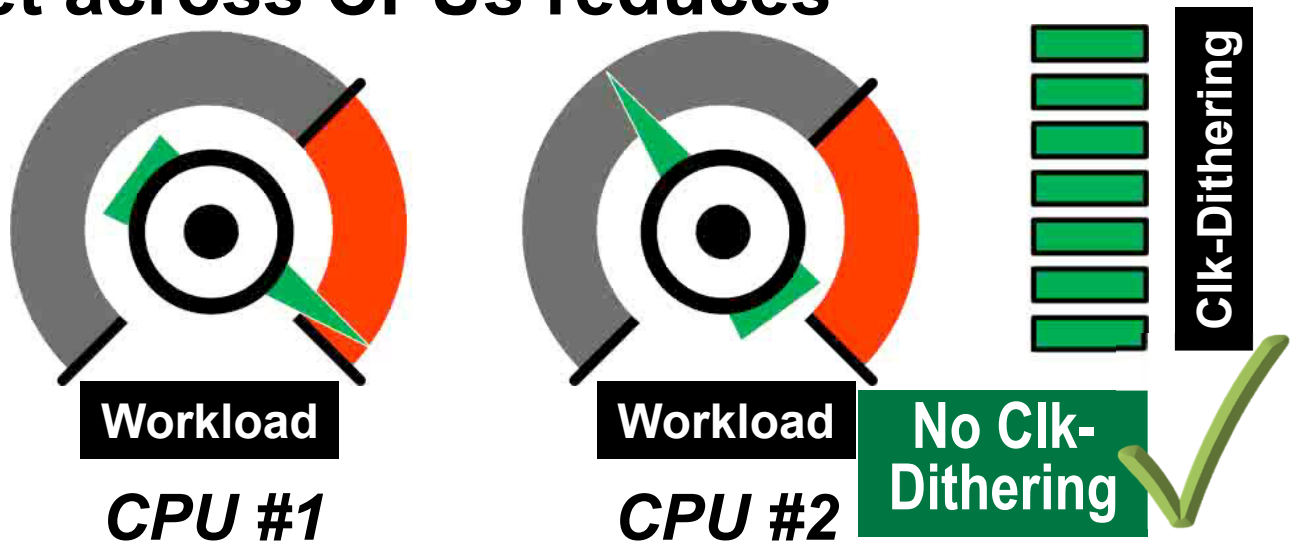

Adaptive Power Allocation (APA)

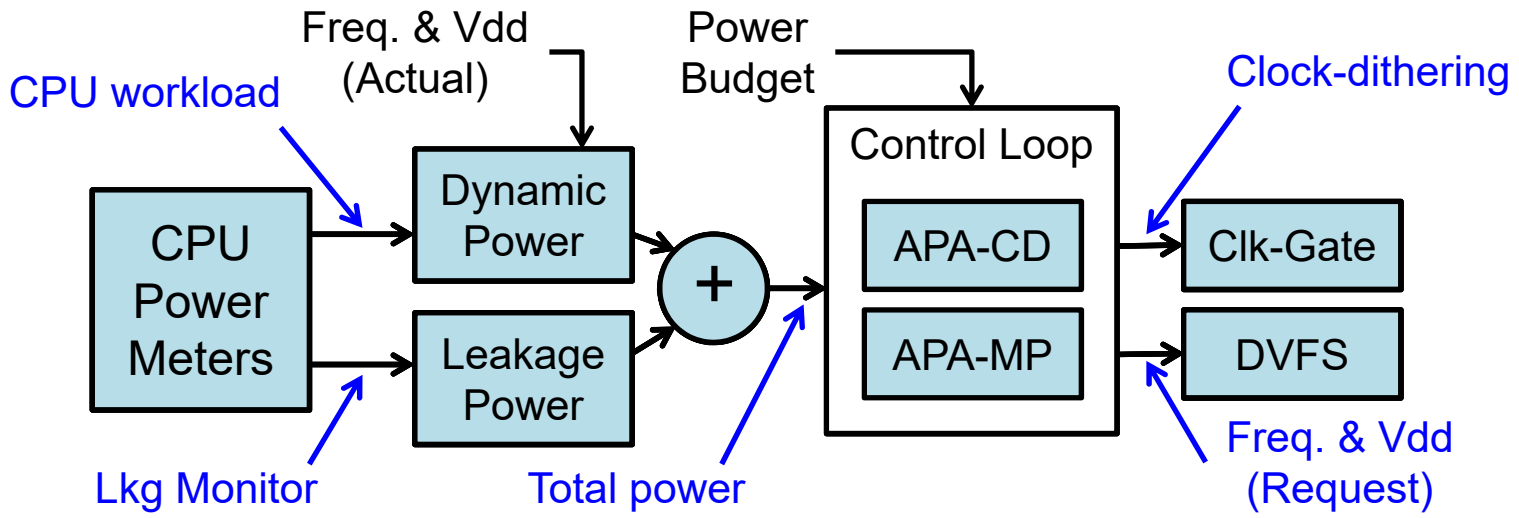

Simulation

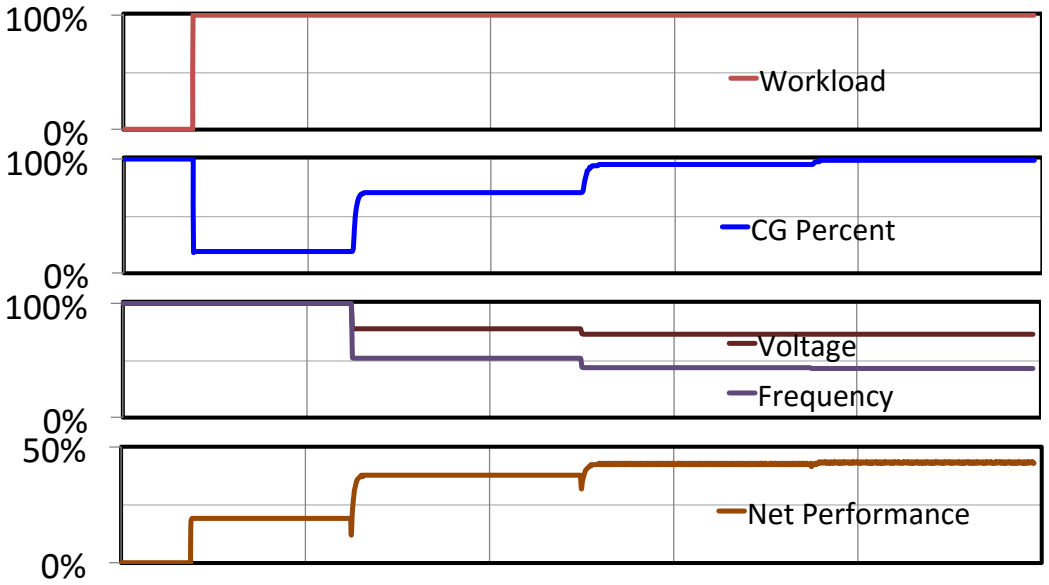

Silicon Measurements

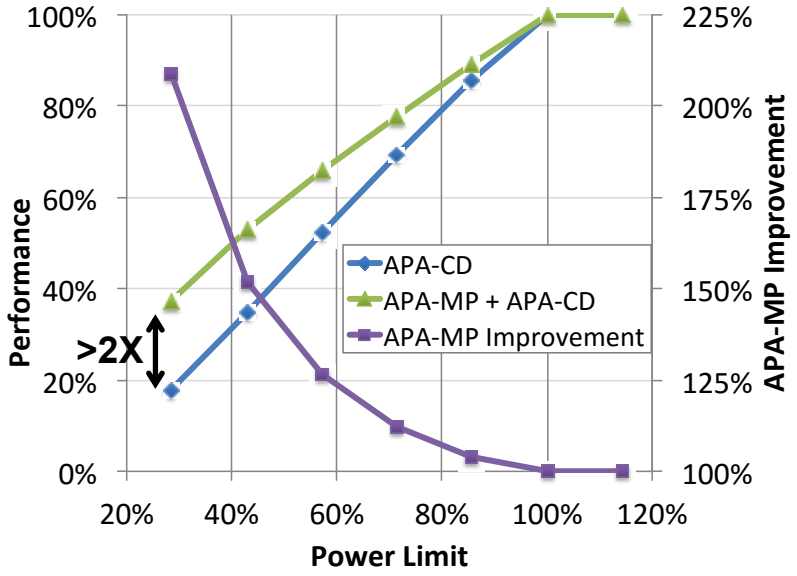

Dynamic Power Meter

- **Event-Accumulator estimation:**
  - Estimation based on Block-Level Events
  - Good correlation at high-power ✓
  - Poor correlation at low-power ✗

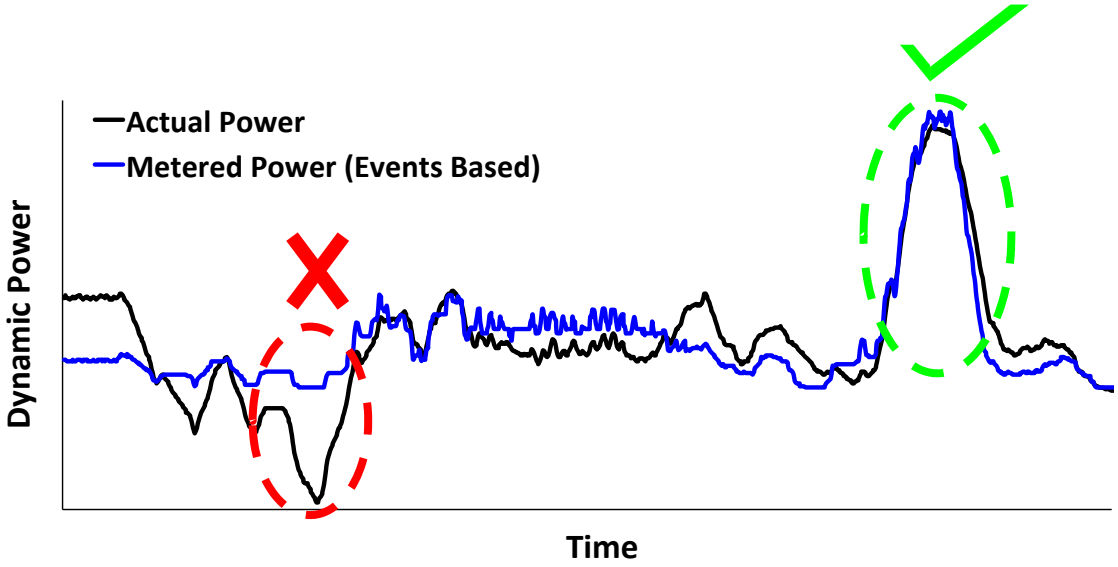

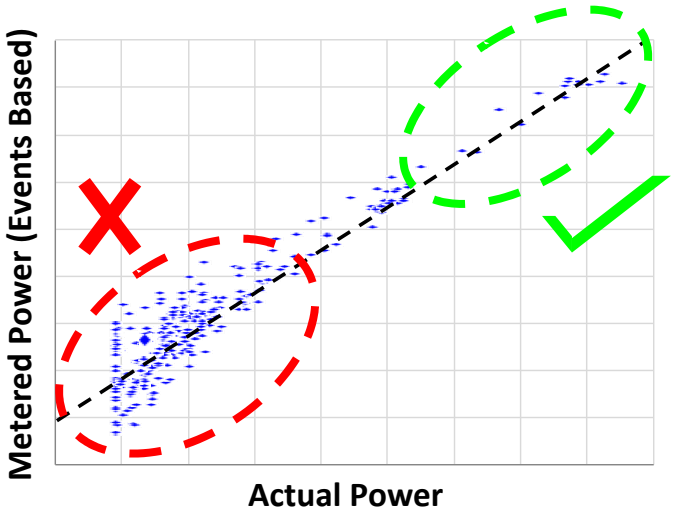

Dynamic Power Meter

- **Clock-Activity-Adder (CAA) estimation:**
  - Estimation based on register clock activity
  - Directly linked to logical clock structure and physical implementation
  - Better correlation vs. Event-based estimation
  - Improved correlation at low-power ✓

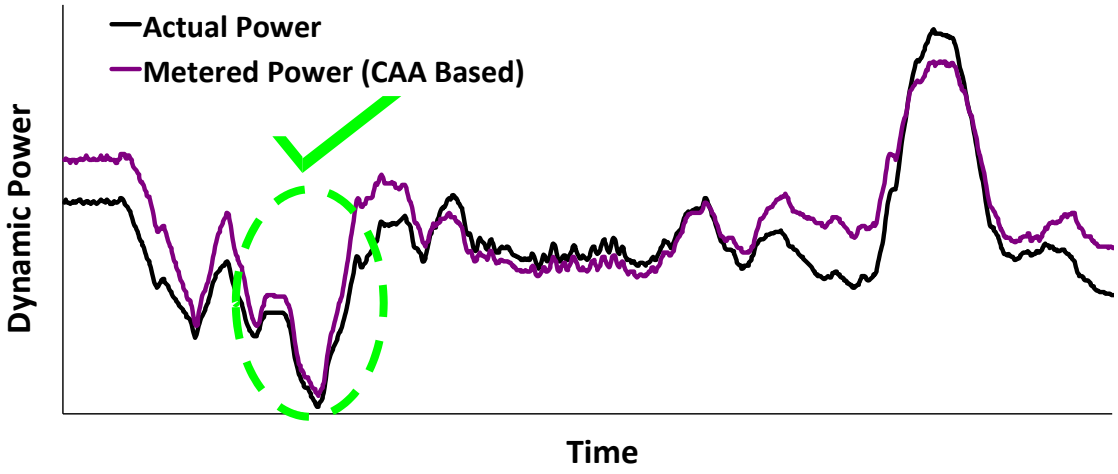

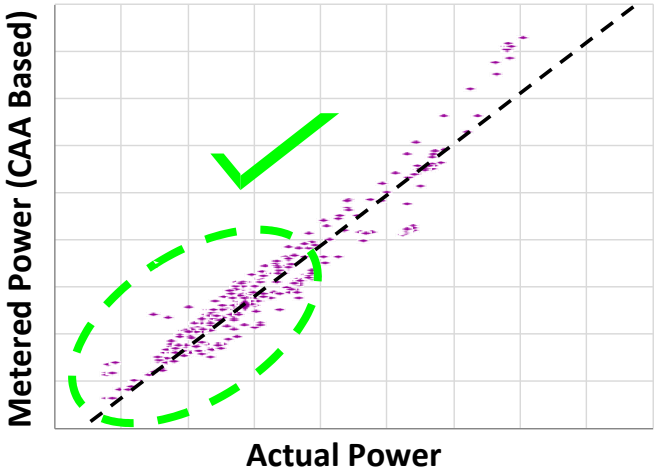

Dynamic Power Meter

- (CAA + Event-Accumulator) estimation:
  - Weighted sum of Events and Clock Activity
  - Comprehends both Block-Level Events and Logic/Physical implementation
  - Improves Correlation further ✓

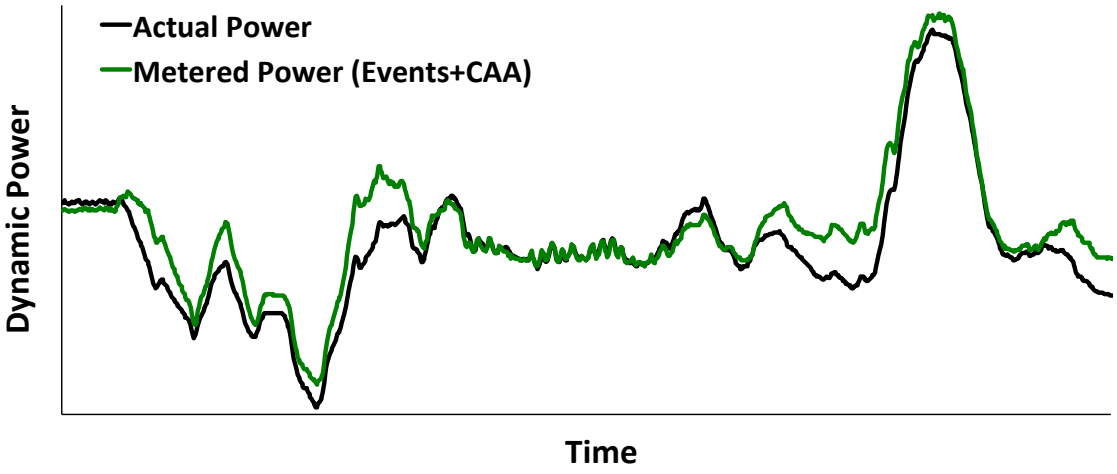

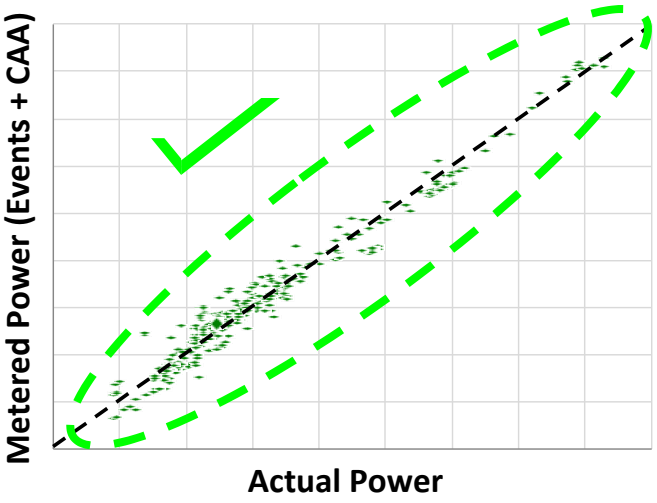

CAA: Implementation

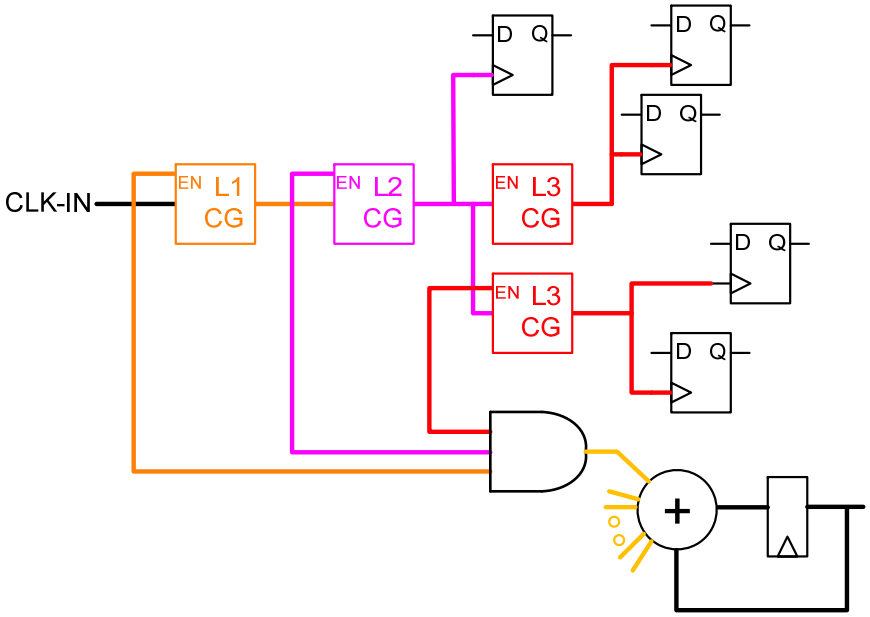

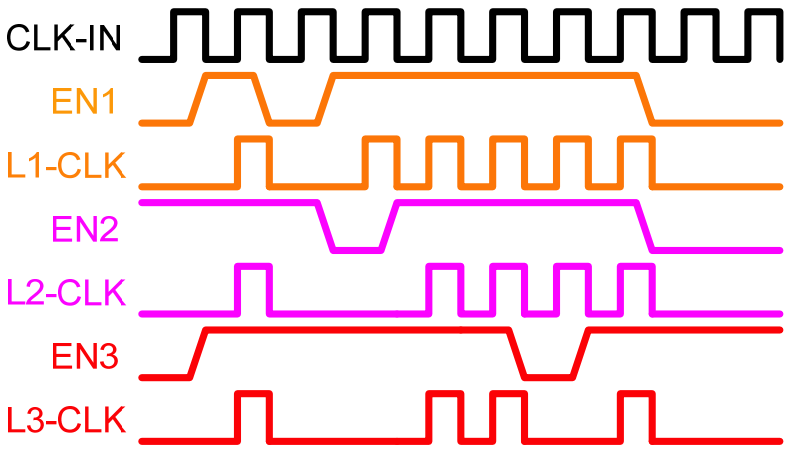

- **Clock-Gate (CG) enable indicates register activity, but requires 'ANDing' of CG enables through the hierarchy**

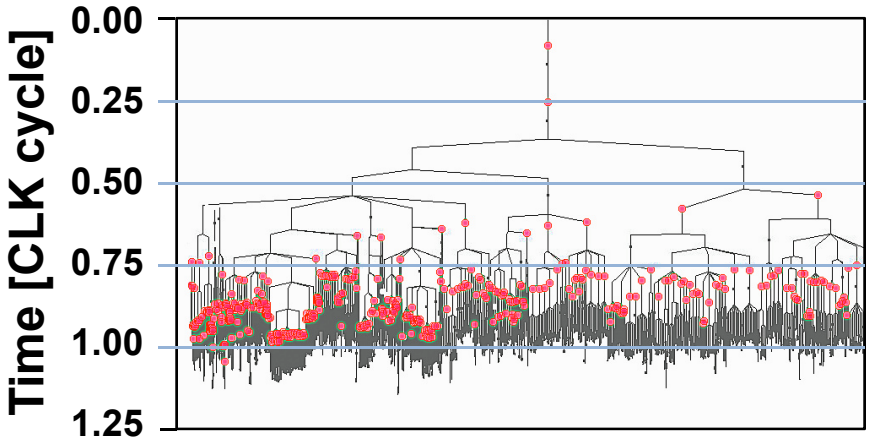

CAA: Implementation

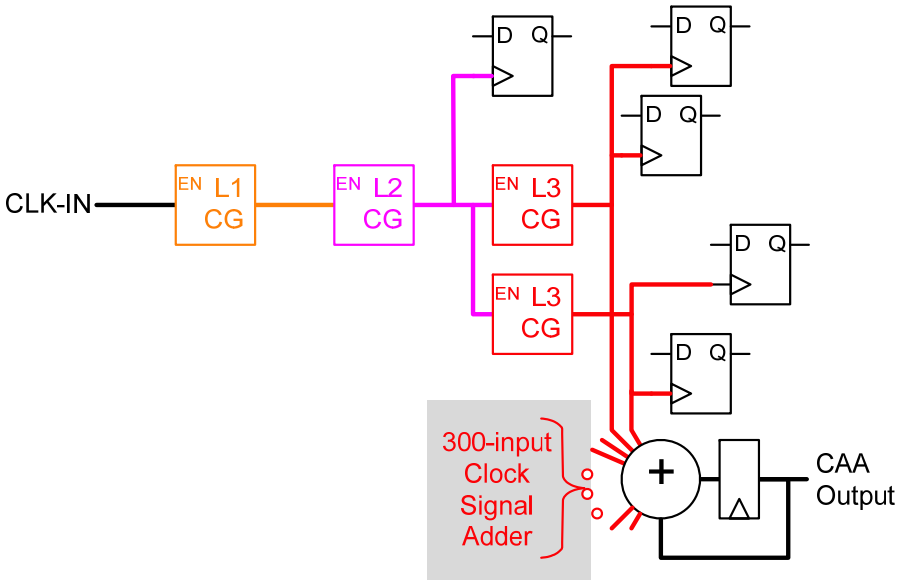

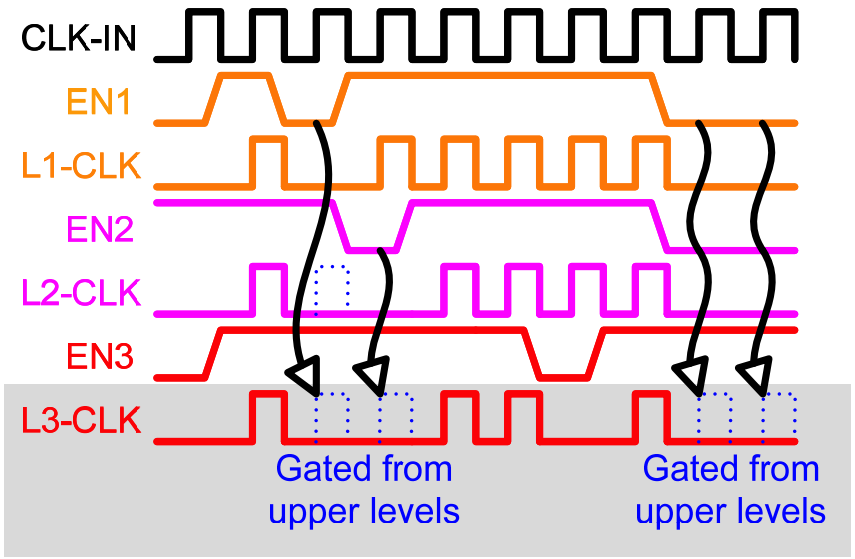

- **Clock-Gate (CG) enable indicates register activity, but requires 'ANDing' of CG enables through the hierarchy**
- **Monitoring clocks directly automatically captures hierarchy**
  - **Simplify logic implementation; immune to clock latency**
- **CG Outputs across clock hierarchy are sampled, summed and accumulated**

Leakage Monitor: Design

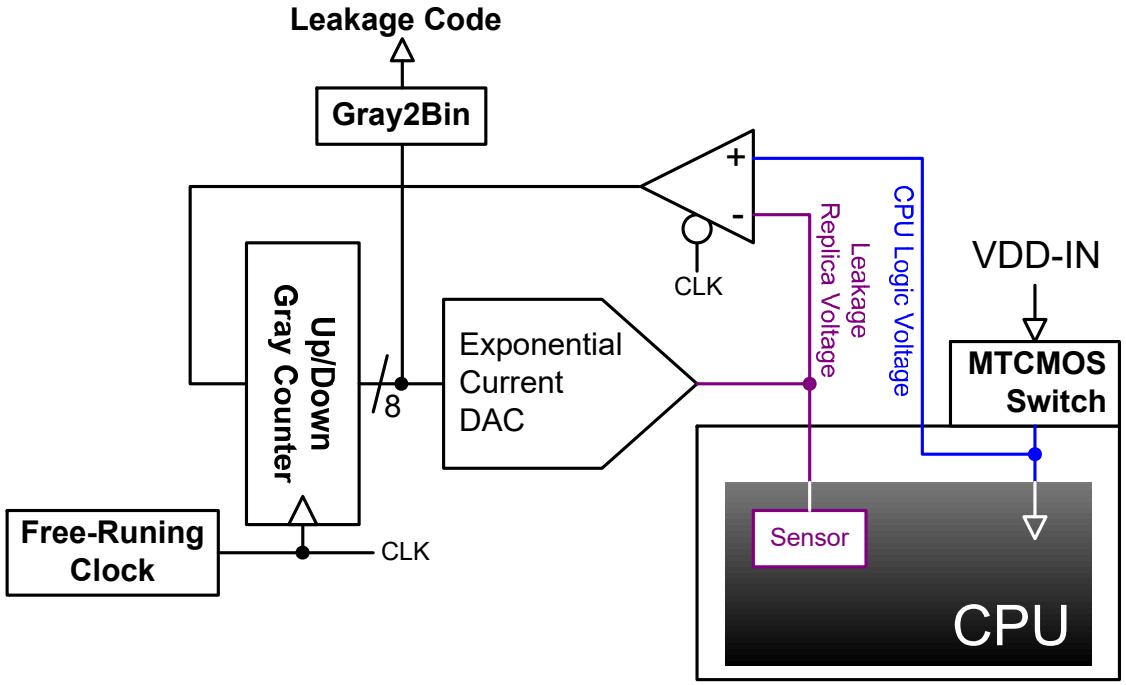

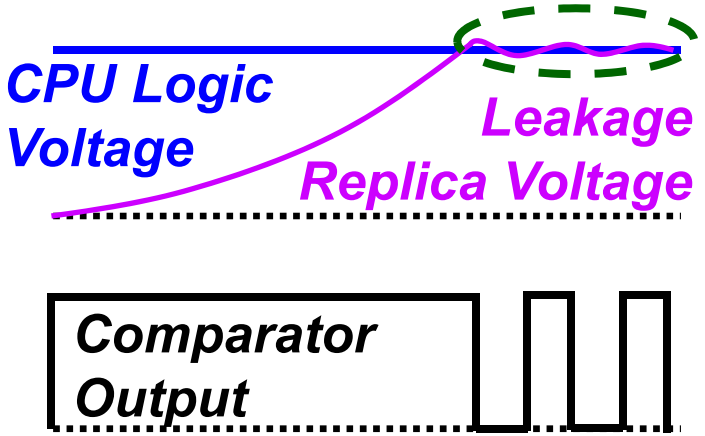

- **Leakage is significant portion of total power**
  - **Varies with Voltage/Temperature; requires constant metering**
- **"Sensor": Similar mix of 'Vt' and 'L' used in CPU**
- **Sensor driven to CPU Voltage by a closed loop:**
  - **Current DAC, Voltage Comparator, Up/Down Counter**
- **Once converged, Counter proportional to leakage current**

Current DAC: Leakage Current Coding

	Normalized Leakage		
	Min	Nom	Max
Process	1	4	26
Temp	1	11	230
Voltage	1	2	3
P-T-V	1	41	5410

$Q < 5\%$

Full-Scale

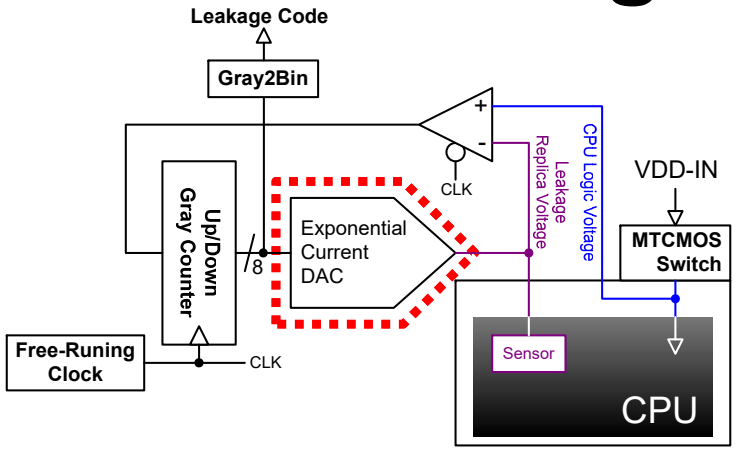

- **#Bits > LOG_2[(*Full-Scale*) / (2**Q*) + 1]**
  - ***Q* is Quantization Error; Target *Q* < 5%**
- **Min : Max P-T-V *Full-Scale* Leakage = 1 : 5410**
  - **Need >15.7 bits of resolution for *Q* < 5%**

Current DAC: Leakage Current Coding

	Normalized Leakage		
	Min	Nom	Max
Process	0.3	1	7
Temp	0.1	1	22
Voltage	0.5	1	1.8
P-T-V	0.02	1	132

Q < 5%

Full-Scale

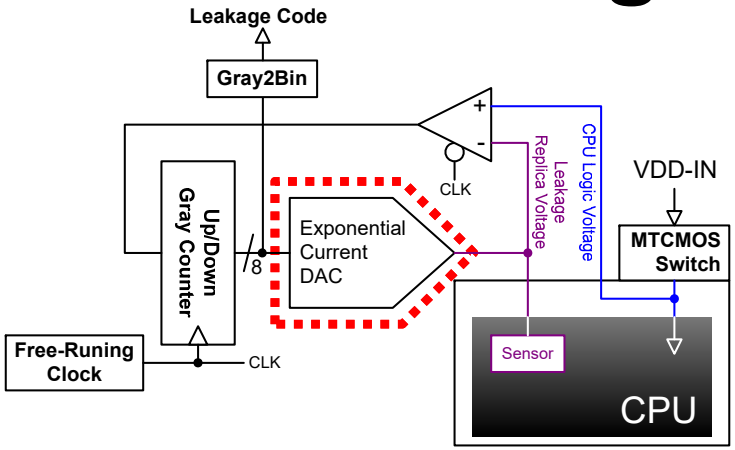

- **#Bits > LOG_2[(*Full-Scale*) / (2**Q*) + 1]**
  - ***Q* is Quantization Error; Target *Q* < 5%**
- **Min : Max P-T-V *Full-Scale* Leakage = 1 : 5410**
  - **Need >15.7 bits of resolution for *Q* < 5%**
- **Nom : Max P-T-V *Full-Scale* Leakage = 1 : 132**
  - **Need >10.3 bits of resolution for *Q* < 5%**

Current DAC: Leakage Current Coding

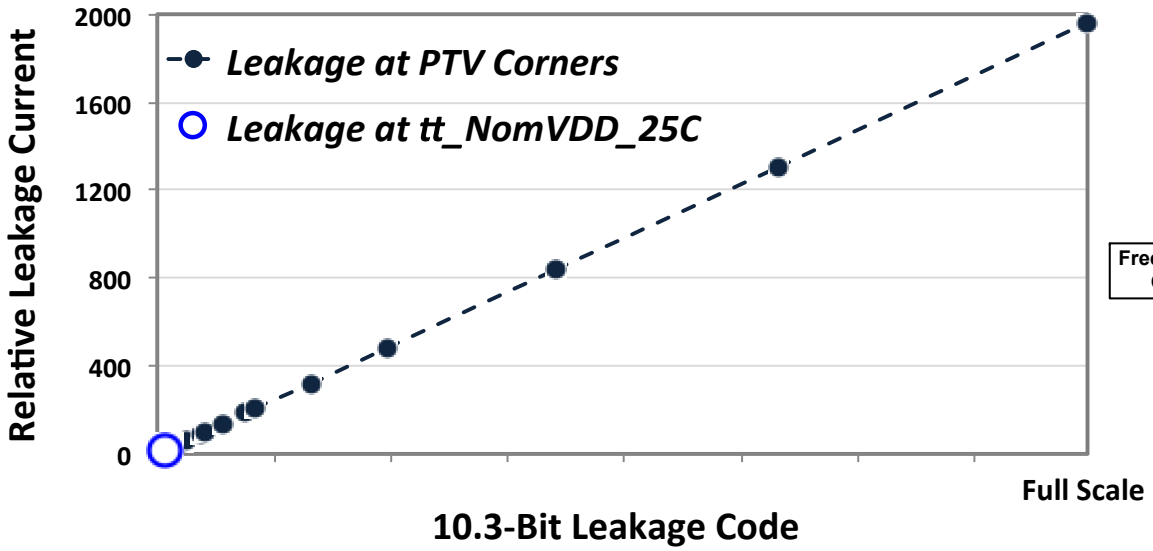

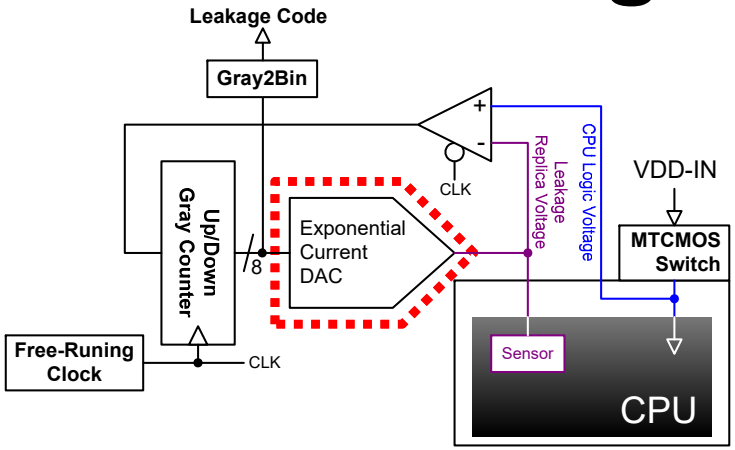

Current DAC: Leakage Current Coding

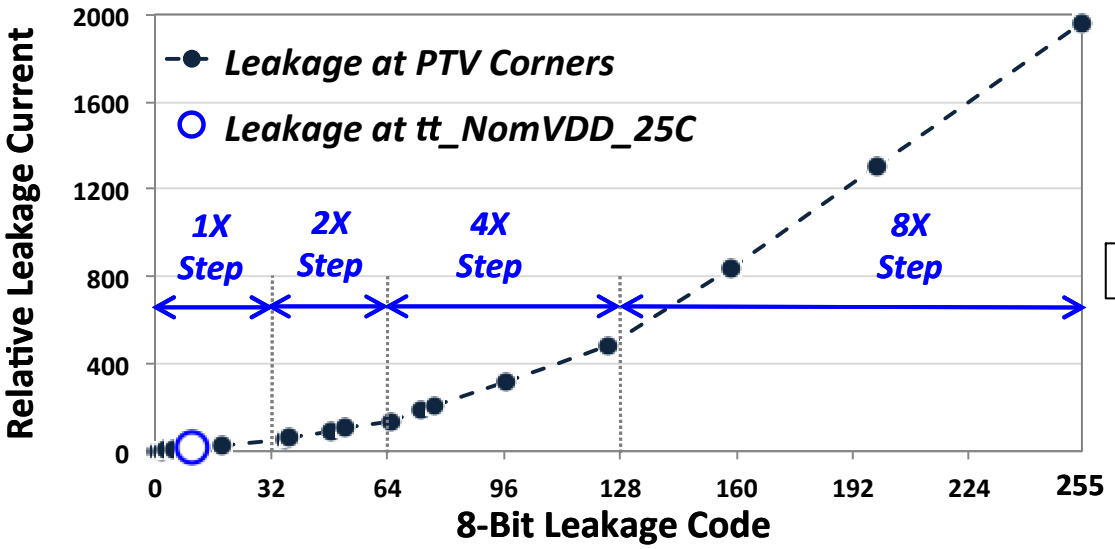

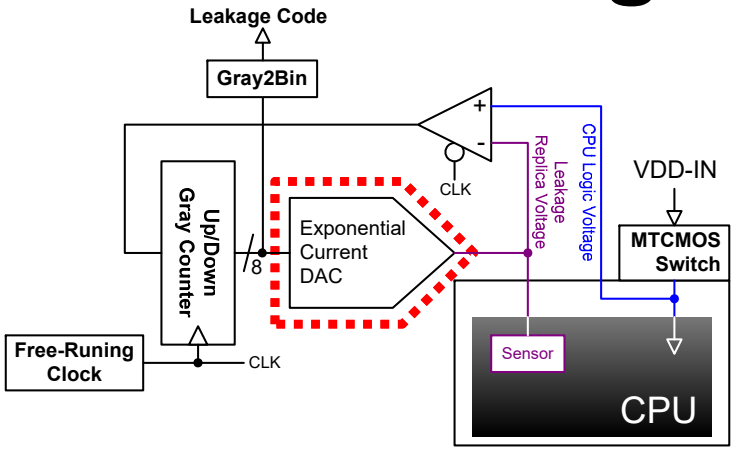

- **Exponential steps to compress 10.3 bits to 8 bits**
  - **Smaller Steps for lower codes**
  - **Exponentially increasing steps for larger codes**
- **10.3 ➔ 8 bits compression reduces complexity**
- ***Q* < 5% from Nom through Max P-T-V**

Exponential Current DAC Schematic

- **Fully segmented Current DAC with 255 switchable current mirrors**
  - **Monotonocity guaranteed**
- **Switchable segments exponentially weighted in 4 groups:**
  - **1X, 2X, 4X, 8X**
- **Total of 1368 1X segments**

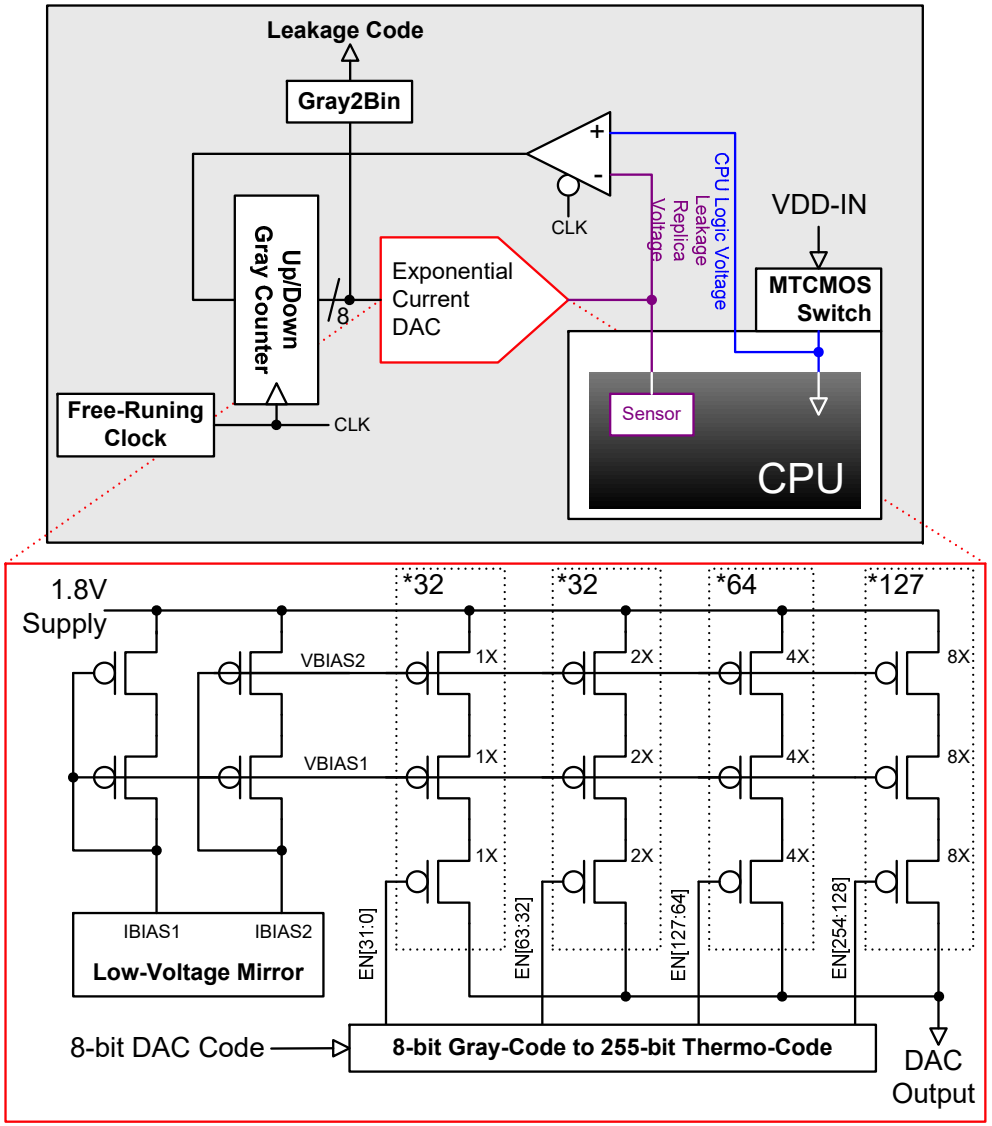

Power Meter: Silicon Correlation

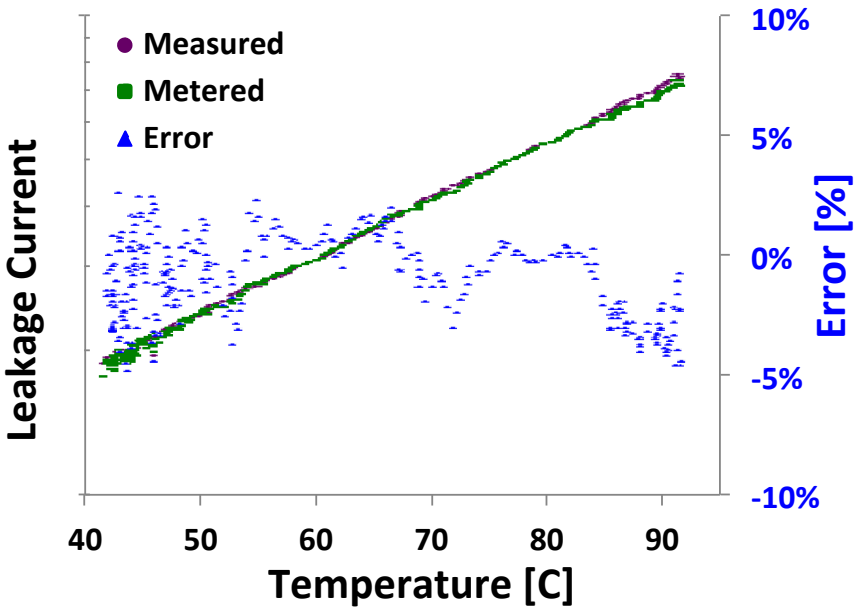

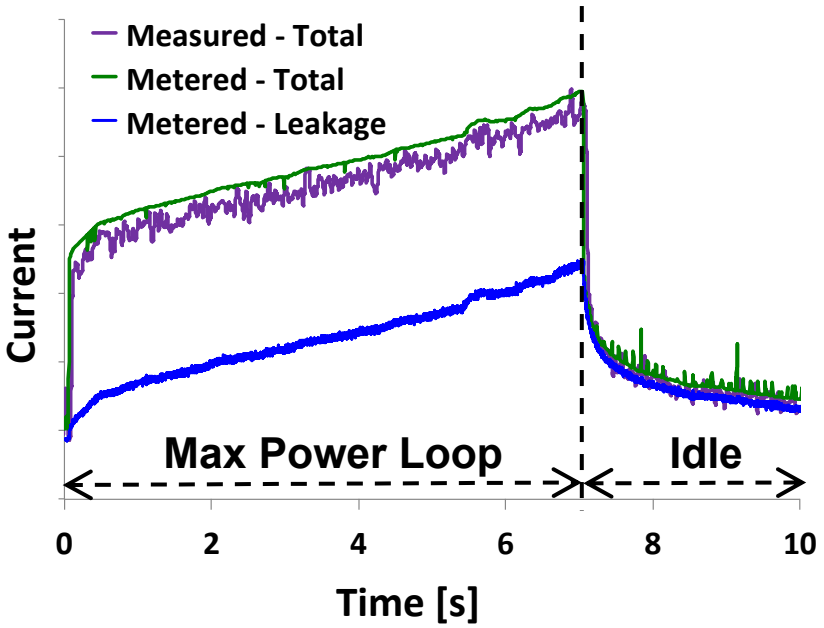

- **Metered Leakage from Leakage Monitor correlates well with Measured Leakage**
  - **Within ±5% across temperature**
- **Good correlation between metered and measured total current**

Outline

- SoC Overview
- Tri-Cluster CPU Subsystem
  - LP Cluster Power-efficiency/Performance benefits
- Adaptive-Power-Allocation (APA)
  - APA Modes: APA-CD and APA-MP
  - Power Meters: Dynamic and Leakage
  - Silicon results
- **SupplEyeScan (SES)**
  - **Motivation**
  - **Design Overview and Operating Modes**
  - **Silicon results**
- Summary

SupplEyeScan (SES): Motivation

- **PDN performance critical to modern day SoCs**
  - **Reduced head-room (VDD – V_T)**
  - **Higher Fmax, current density**
- **Quality PDN: Effort and resource intensive**
  - **PCB-Package-Die Co-design**
  - **Package and PCB components**
- **Need simple/efficient way to validate PDN**
  - **"SupplEyeScan" (SES), a compact, on-die, high-BW oscilloscope was developed**

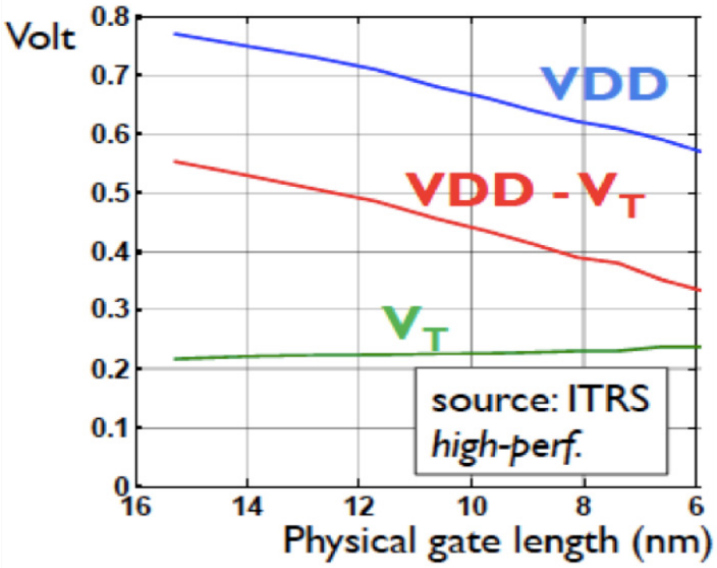

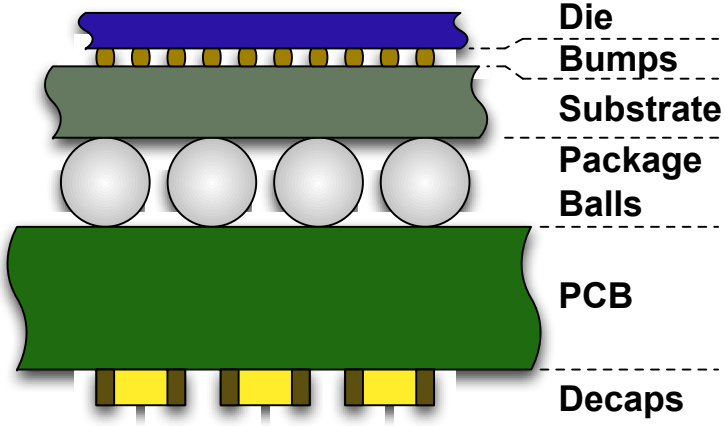

SupplEyeScan: Design

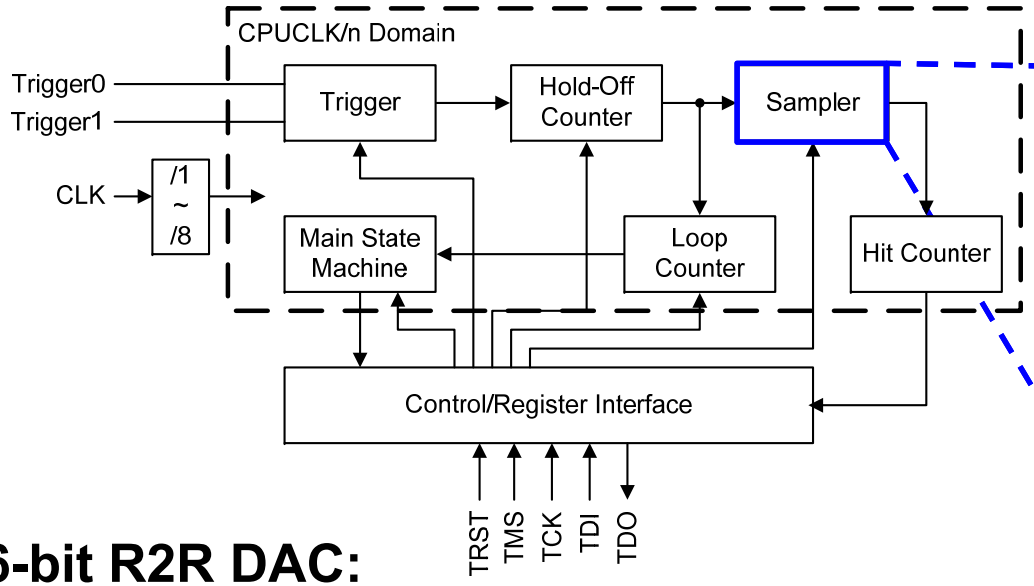

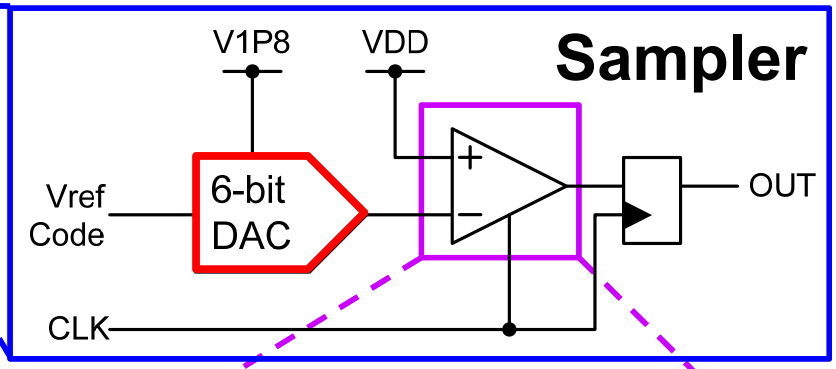

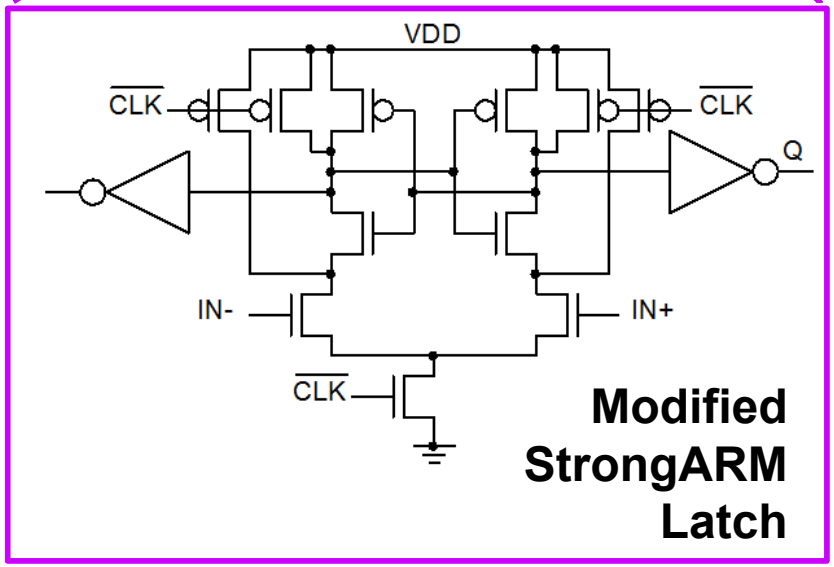

6-bit R2R DAC:

- Output Reference Voltage: 0.6V~1.2V
- DAC Step Size: ~10mV
- Powered by 1.8V Supply

Modified StrongARM latch[5]:

- Compares die-internal supply with reference voltage from DAC
- Operates at CPU clock frequency
- >1GHz Measurement BW

[5] Y. T. Wang and B. Razavi, "An 8-bit 150-MHz CMOS A/D Converter," IEEE J. Solid-State Circuits, vol. 35, pp. 308–317, Mar. 2000.

SES: Histogram Mode

- **Histogram Mode captures higher probability portion of the distribution through sub-sampling of VDD**

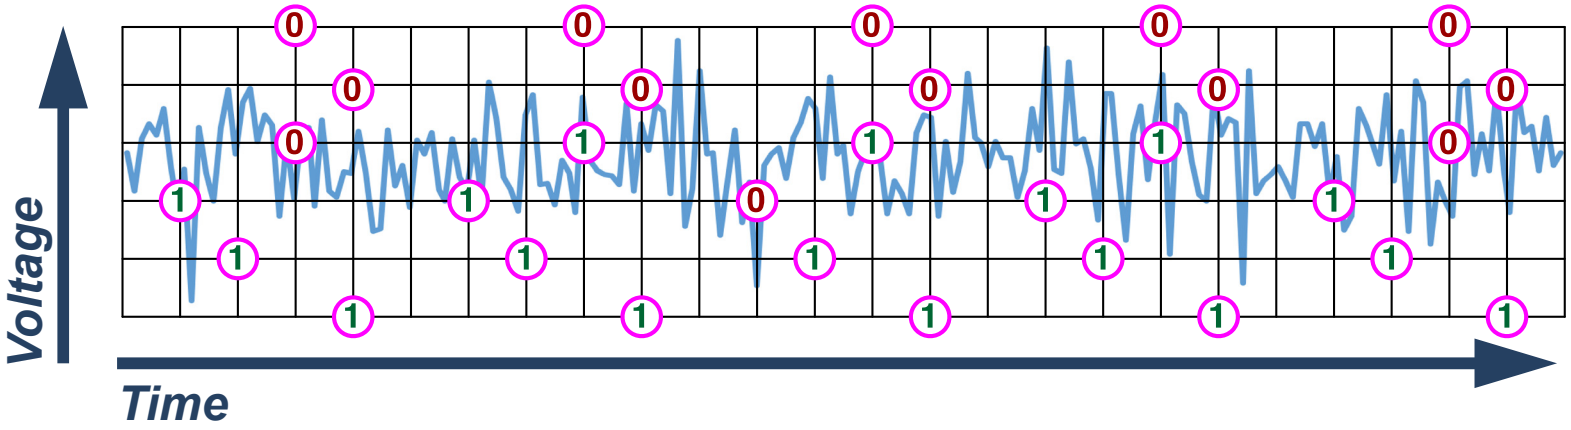

- **Up to 255 samples of (VDD – VREF) captured:**
  - **Spaced to capture low-frequency VDD variation**
- **DAC reference (VREF) swept from min to max**
- **Probability of (VDD>VREF) measured at each VREF**

Histogram Mode: Silicon results

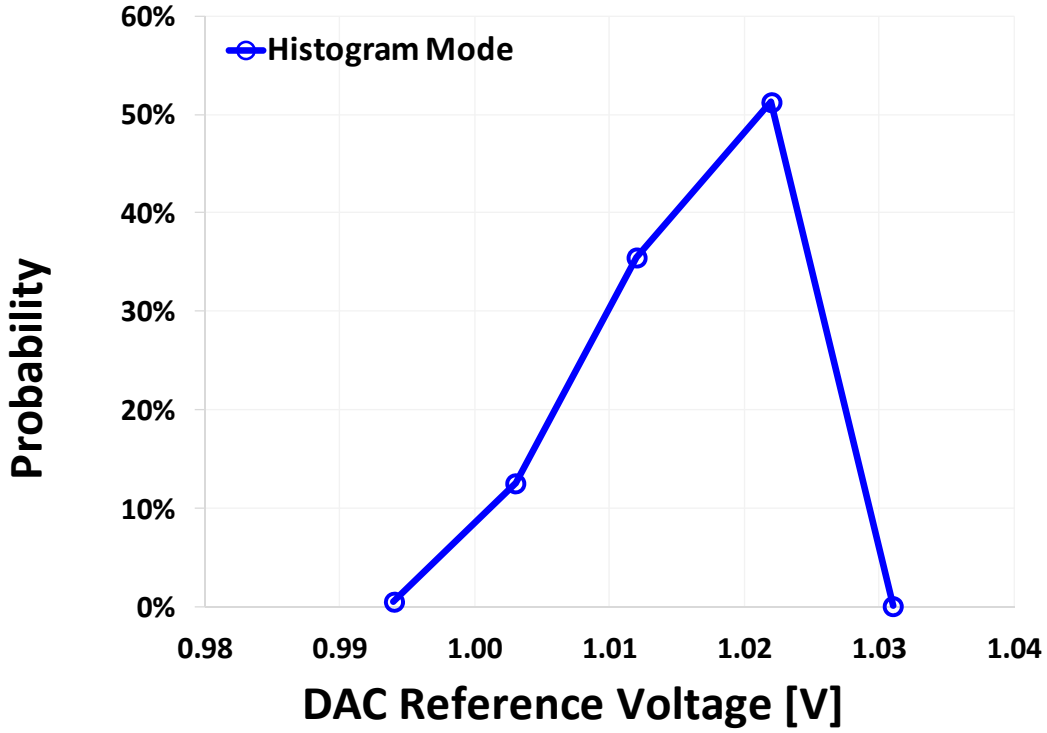

- **Histogram mode captures high probability statistics of supply voltage**
  - **Infrequent voltage excursions not captured**

SupplEyeScan Peak-Detect Mode

- **Peak-Detect Modes capture "tail" of the distribution through continuous-sampling of VDD**

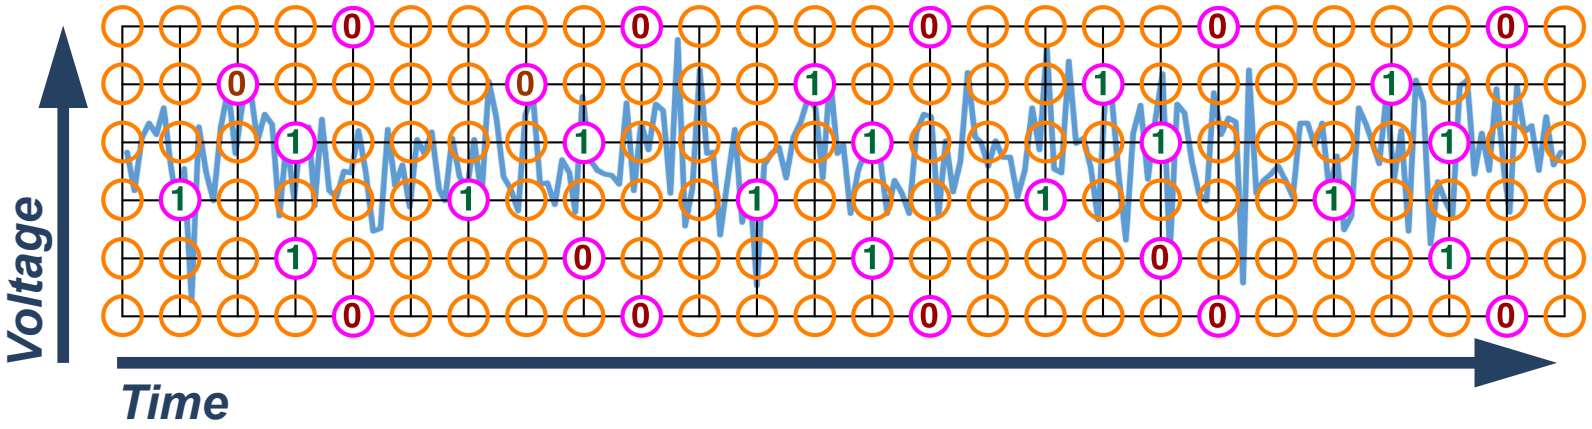

- **Within each window, record any occurrence of:**
  - **Peak-Low: VDD < VREF**
  - **Peak-High: VDD > VREF**
- **Capture worst-case min/max voltages**

Peak-Detect Mode: Silicon results

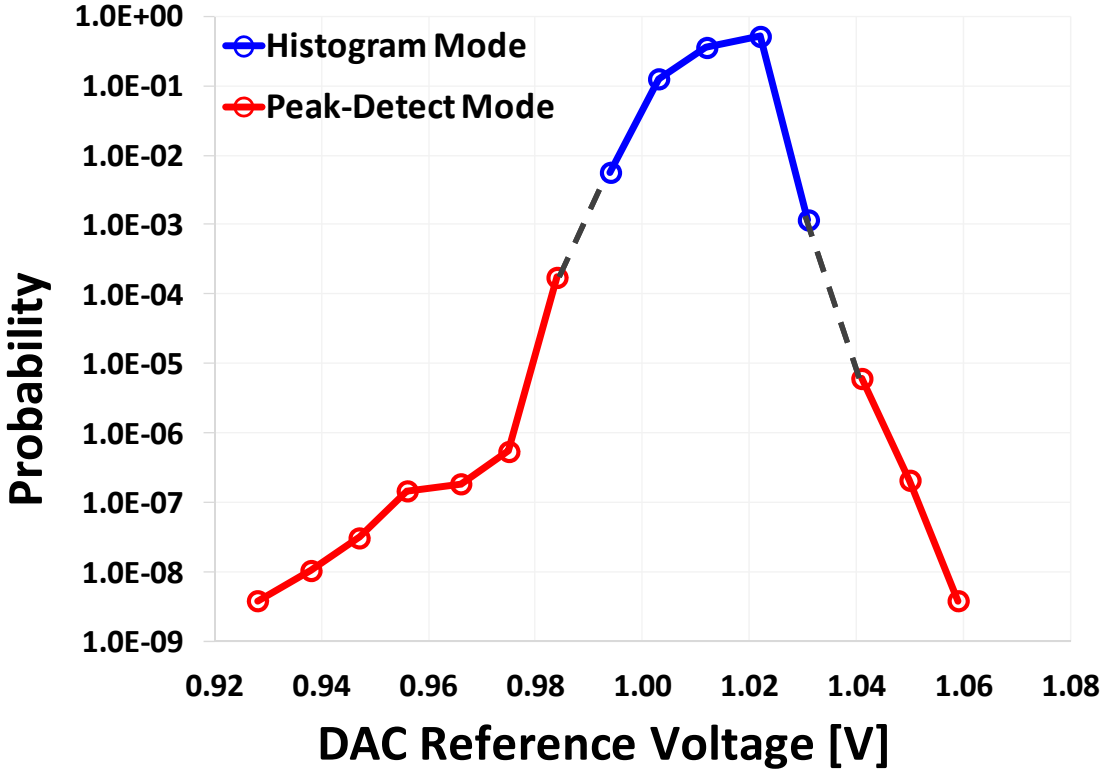

- **Peak-Detect (high and low) modes capture "tail" statistics of supply voltage**
- **Histogram and Peak-Detect modes combined provide overall PDN quality**

SupplEyeScan Transient Mode

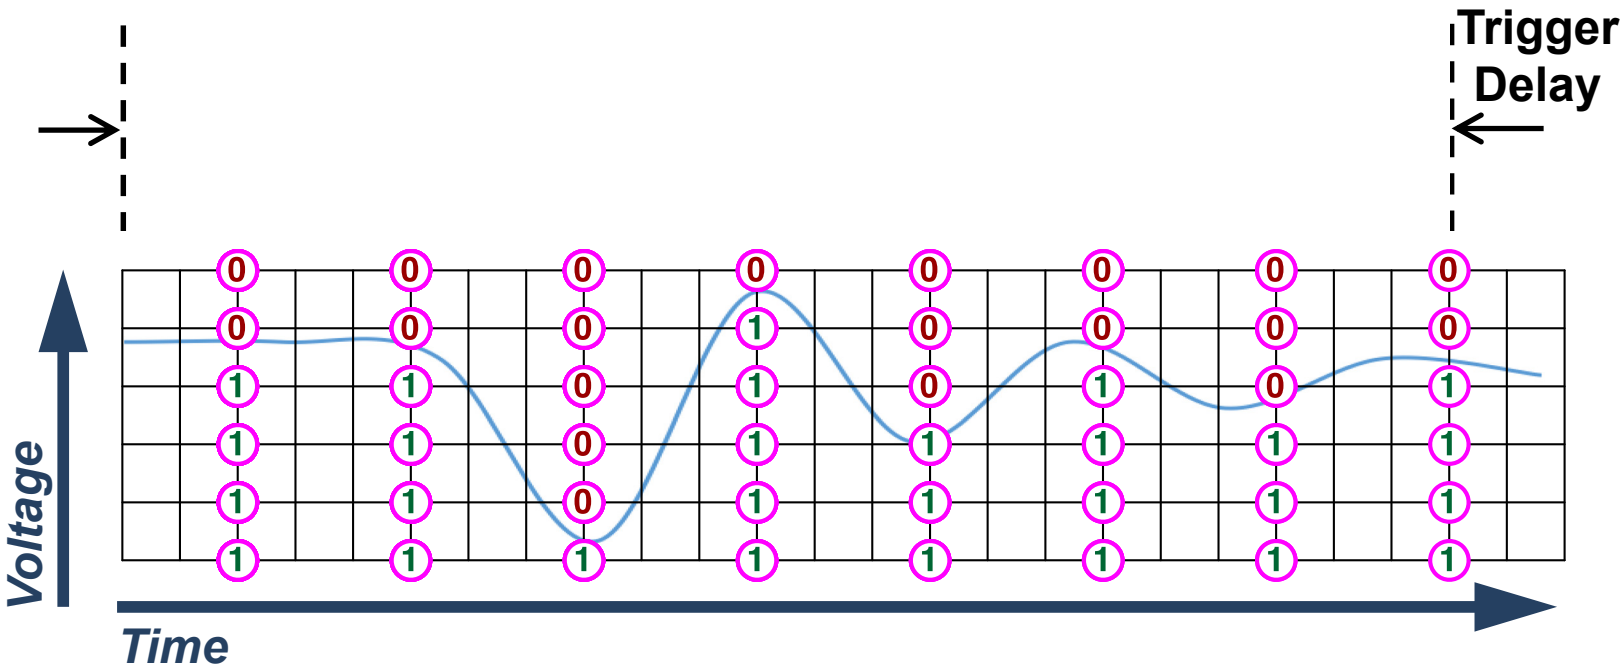

- Software sequence designed to stress dI/dT
  - Issues a hardware trigger synchronous to dI/dT event (SEV instruction)
- Sweep VREF and Trigger delay to capture transient waveform

Transient Mode: Silicon Results

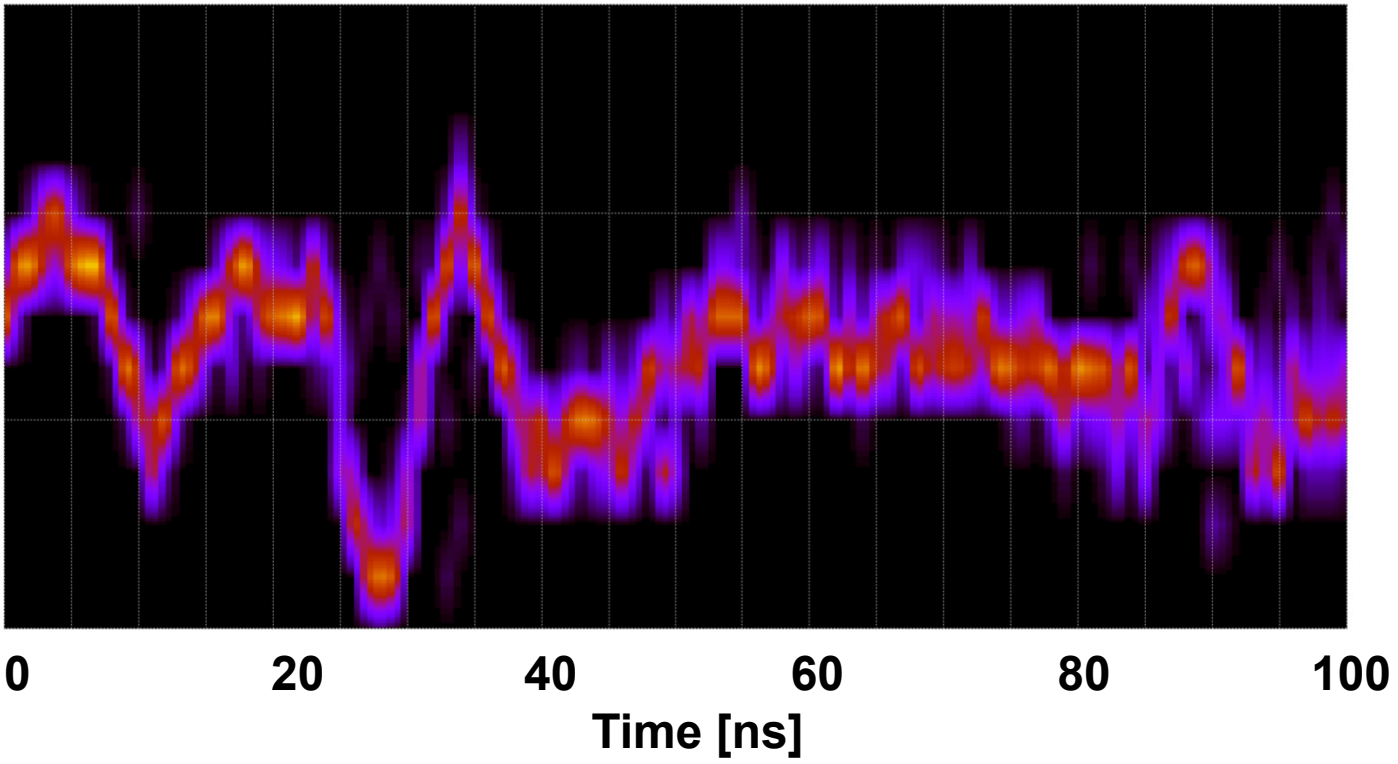

- **SupplEyeScan transient mode waveform**
  - **Each Pixel is (Trigger Delay, VREF) co-ordinate**
  - **Probability of VDD = VREF shown on color axis at each pixel**

Outline

- SoC Overview
- Tri-Cluster CPU Subsystem
  - LP Cluster Power-efficiency/Performance benefits
- Adaptive-Power-Allocation (APA)
  - APA Modes: APA-CD and APA-MP
  - Power Meters: Dynamic and Leakage
  - Silicon results
- SupplEyeScan (SES)
  - Motivation
  - Design Overview and Operating Modes
  - Silicon results
- **Summary**

Summary

- **A 20nm 2.5GHz Tri-Cluster CPU Sub-system:**
  - ✓ New LP Cluster bridges power-efficiency/performance gap between ULP and HP Clusters
  - ✓ +40% performance vs. ULP
  - ✓ +40% power efficiency vs. HP
- **Adaptive Power Allocation (APA):**
  - ✓ 2X Performance when power limited
  - ✓ Power Meters (Dynamic and Leakage)
- **SupplEyeScan for Simple/Efficient PDN validation:**
  - ✓ Supply Voltage Variation (Histogram, Peak-Detect Mode)
  - ✓ Time domain dI/dT response (Transient Mode)

Thank You

A 197mW 70ms-Latency Full-HD 12-Channel Video-Processing SoC for Car Information Systems

S. Mochizuki, K. Matsubara, K. Matsumoto,
C. Nguyen*, T. Shibayama, K. Iwata, K. Mizumoto,
T. Irita**, H. Hara**, T. Hattori

Renesas System Design, Tokyo, Japan
* Renesas Design Vietnam, HCMC, Vietnam
** Renesas Electronics, Tokyo, Japan

Agenda

- Background
- Chip specifications and video performance
- Overview of the key technologies
- Memory bandwidth reduction
- Experimental results
- Conclusion

Integrated cockpit systems (ICS)

- Future car information systems will support
  - both infotainment and driver assistance (ADAS)

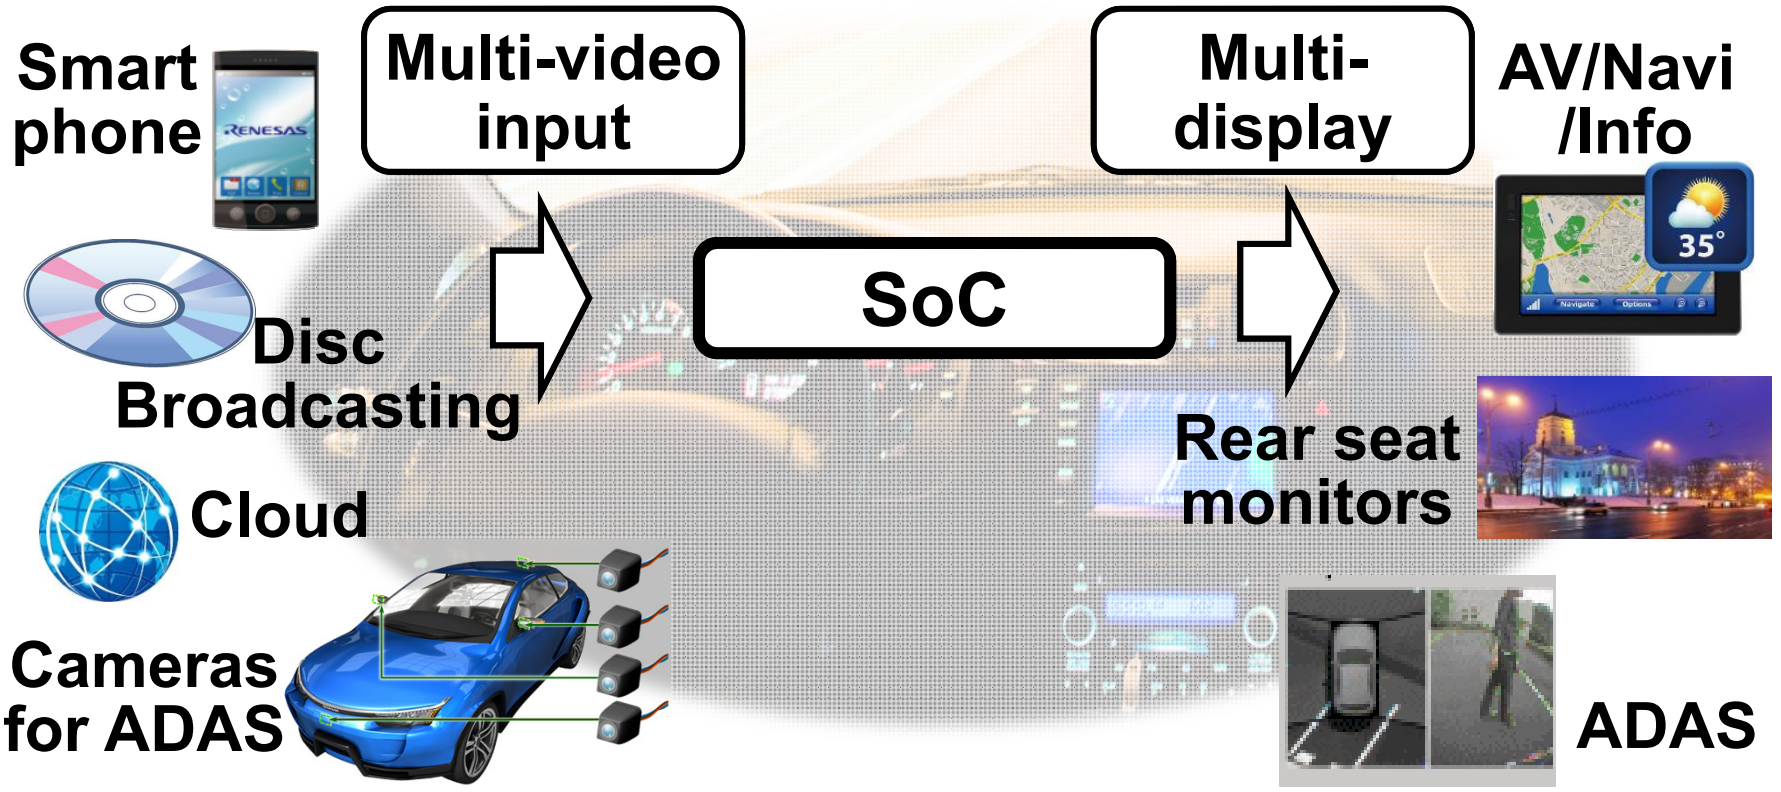

Integrated cockpit systems (ICS)

- Future car information systems will support
  - both infotainment and driver assistance (ADAS)
  - in-car video transfer via Ethernet

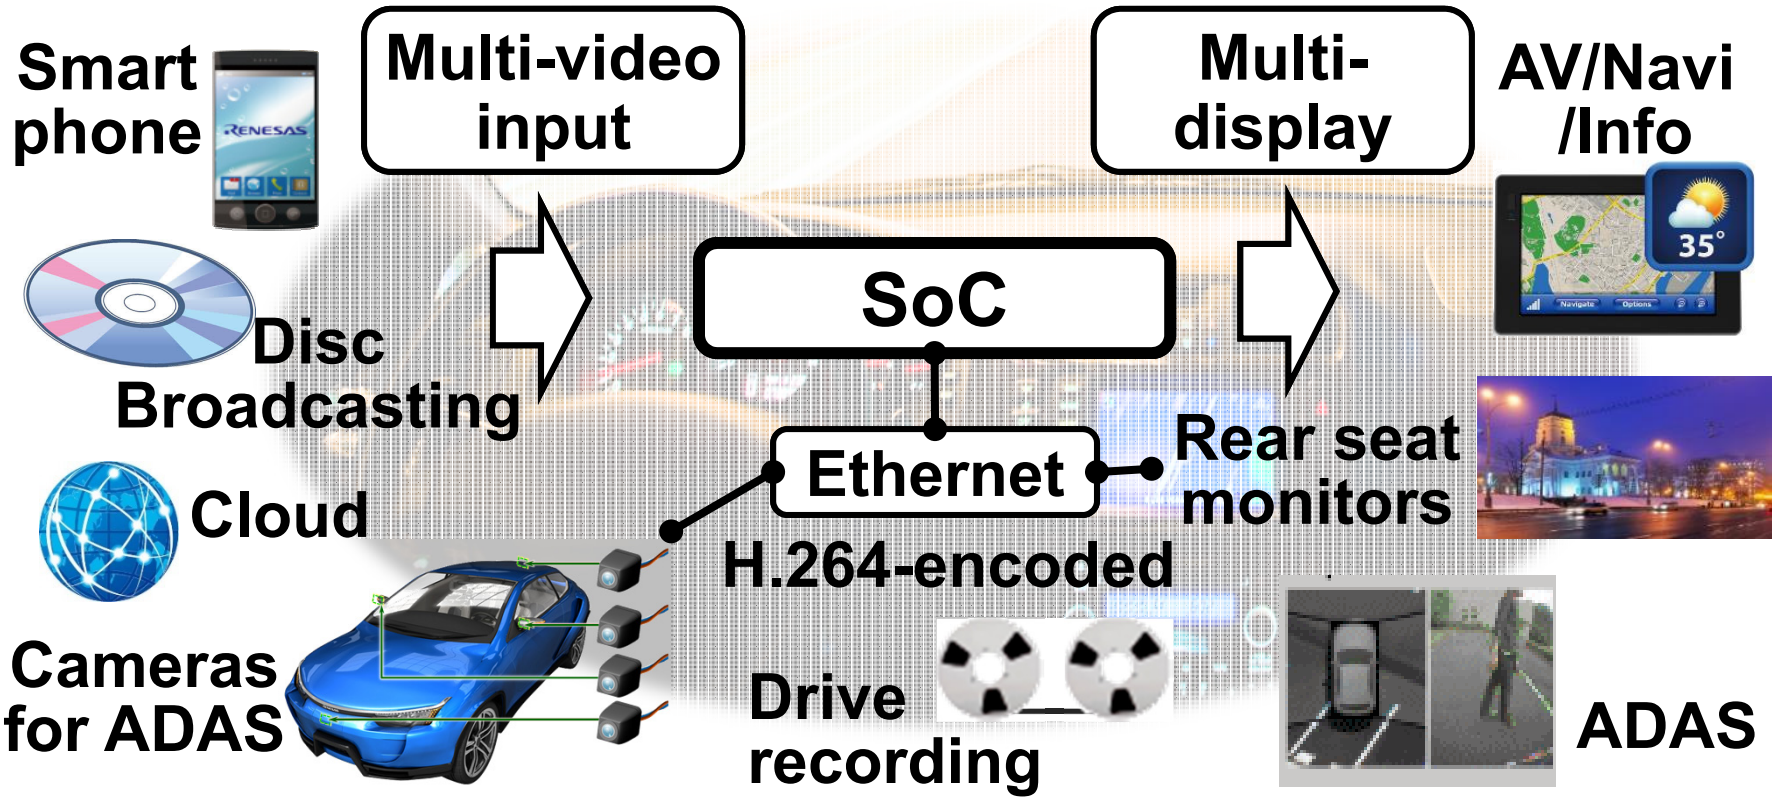

Integrated cockpit systems (ICS)

- **Future car information systems will support**
  - **both infotainment and driver assistance (ADAS)**
  - **in-car video transfer via Ethernet**

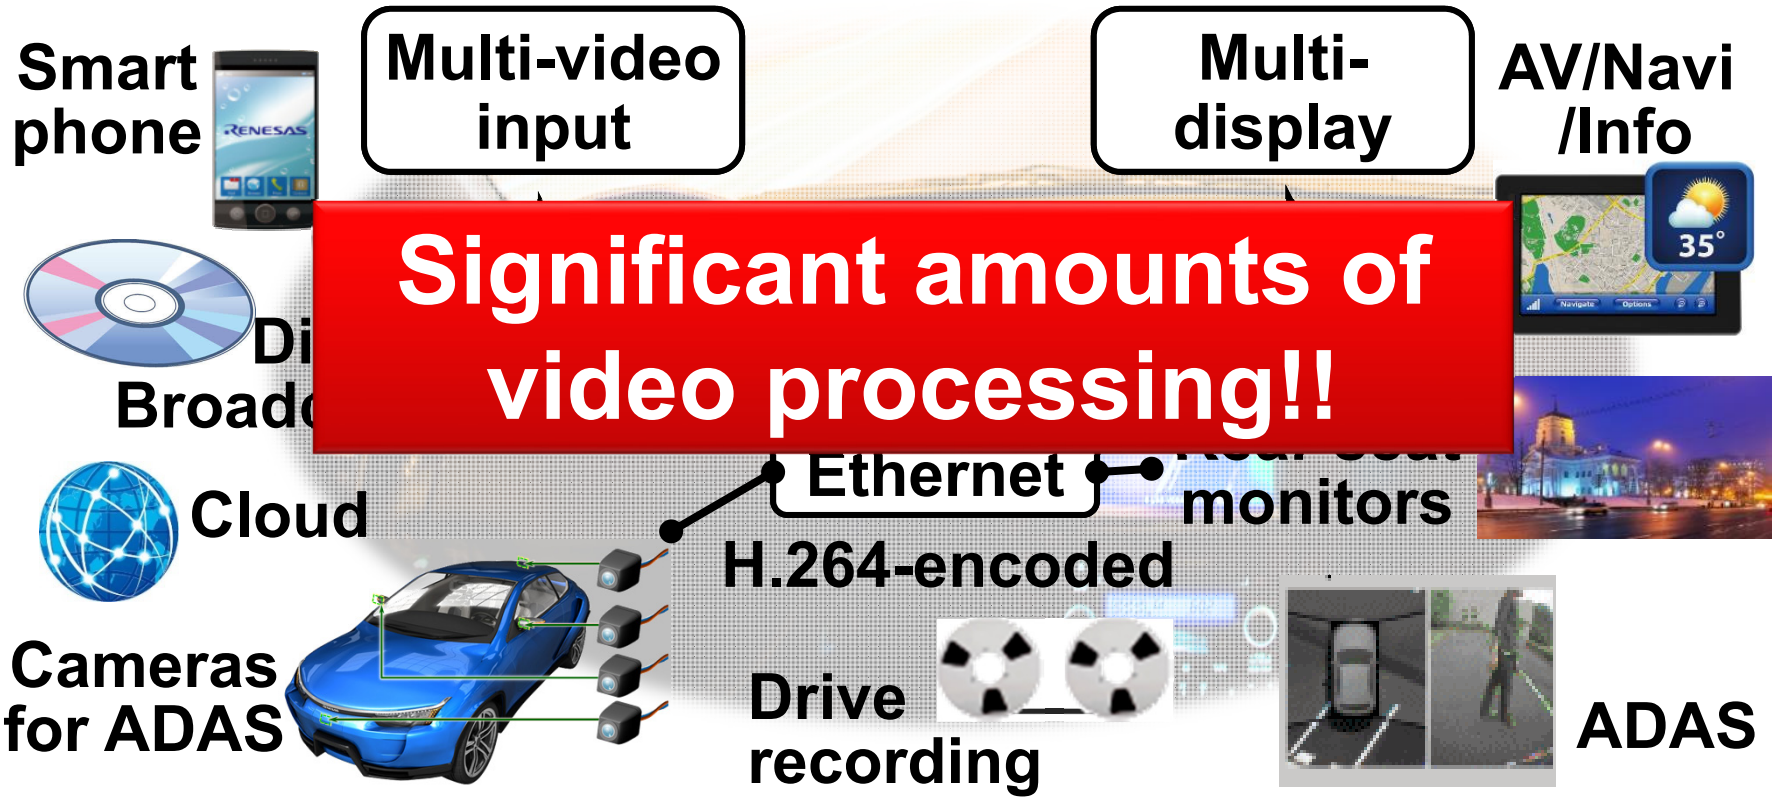

Chip specifications

Technology		16nm FinFET CMOS
Chip size		16.4mm^2 (Video)
Supply voltage		0.8V (core)
Clock frequency		400MHz (Video)
External memory		LPDDR4-3200
Video codec	Performance	1920x1080 x 30fps x 12ch @H.264
	Resolution	4096x2304 @H.265/HEVC, H.264
	Standard	H.265/HEVC, H.264, MPEG-2/4, VC-1, VP8
Video processing		De-interlace, Scaling, Blending, Color adjustment, De-skew, others

Chip micrograph

- The video subsystem in the SoC.

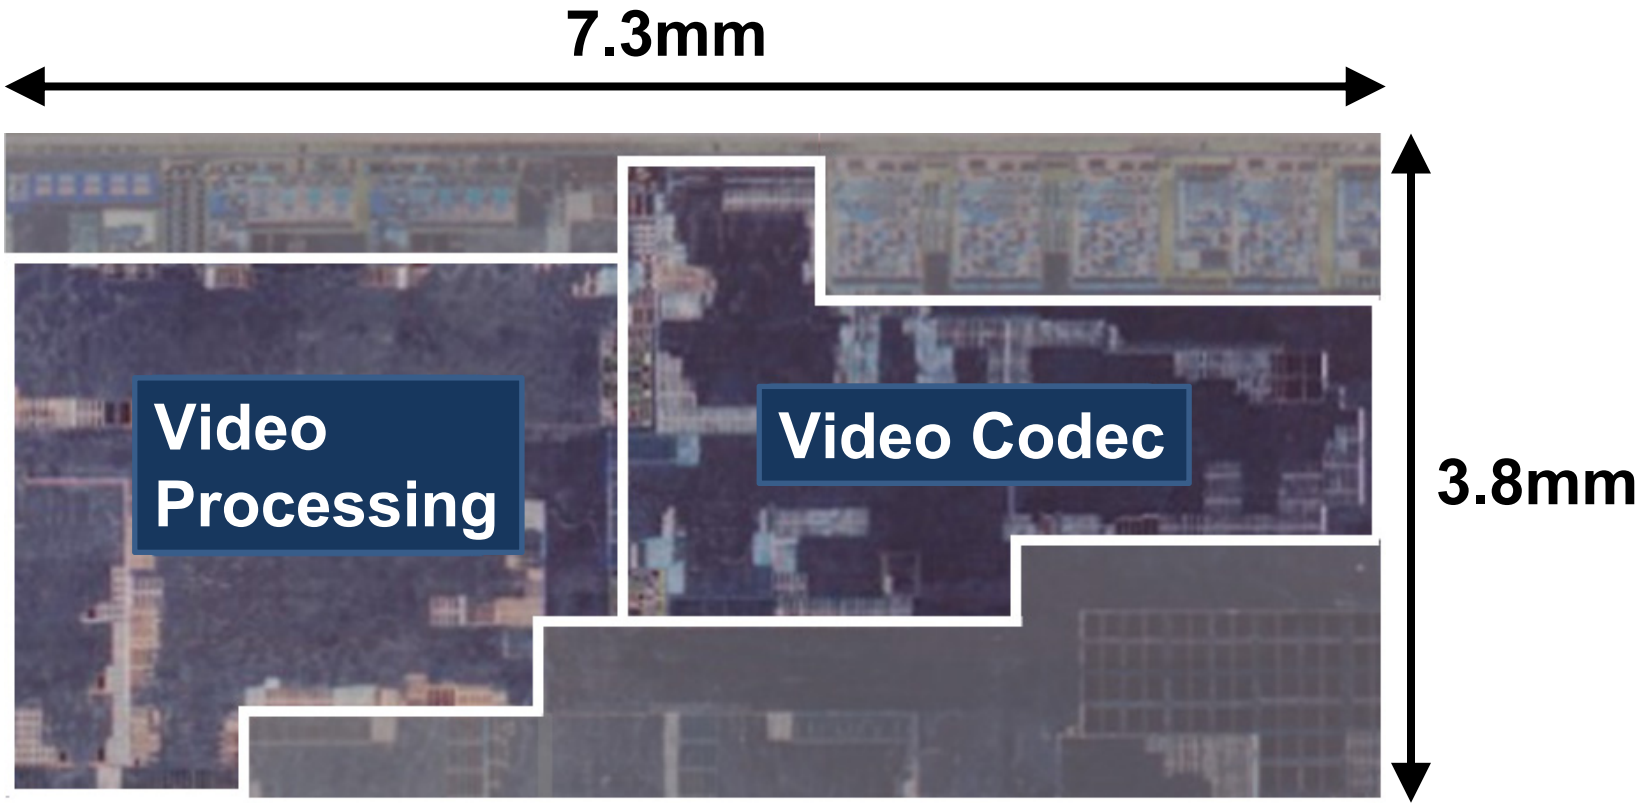

Video performance

- **750Mpixels/s exceeds previous works and a trend for application processors (AP).**

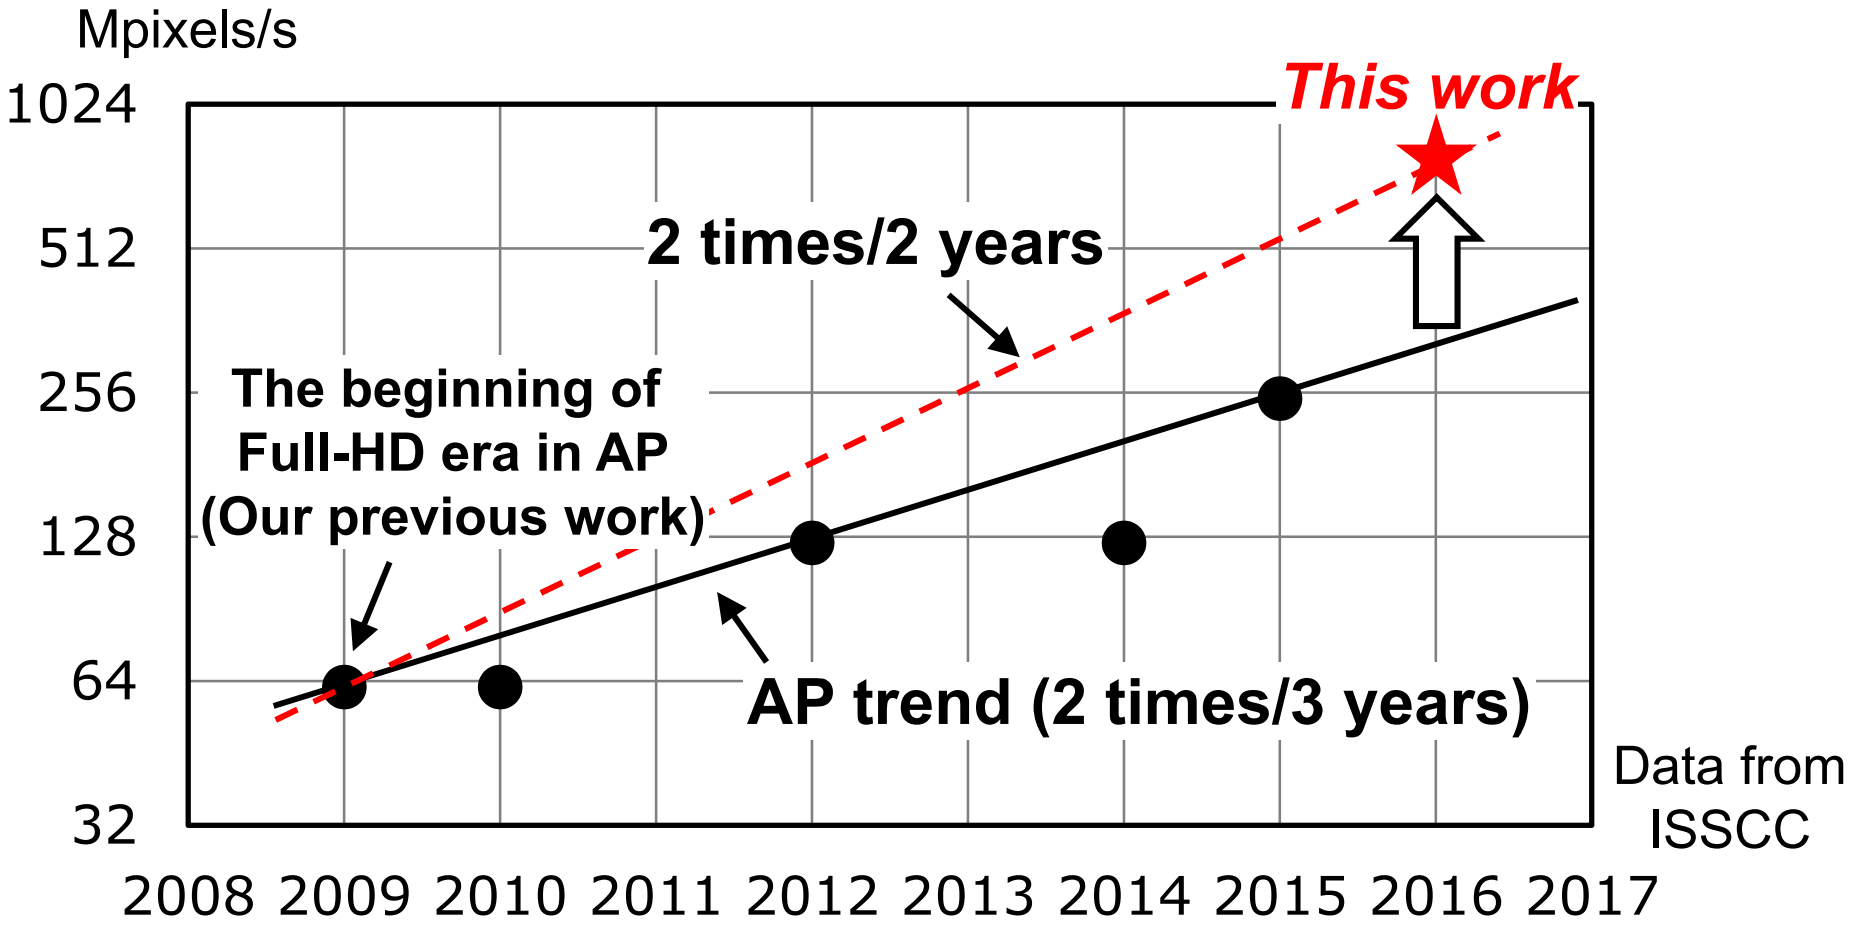

Block diagram

- **17 video processors of 6 different types for infotainment and driver assistance.**

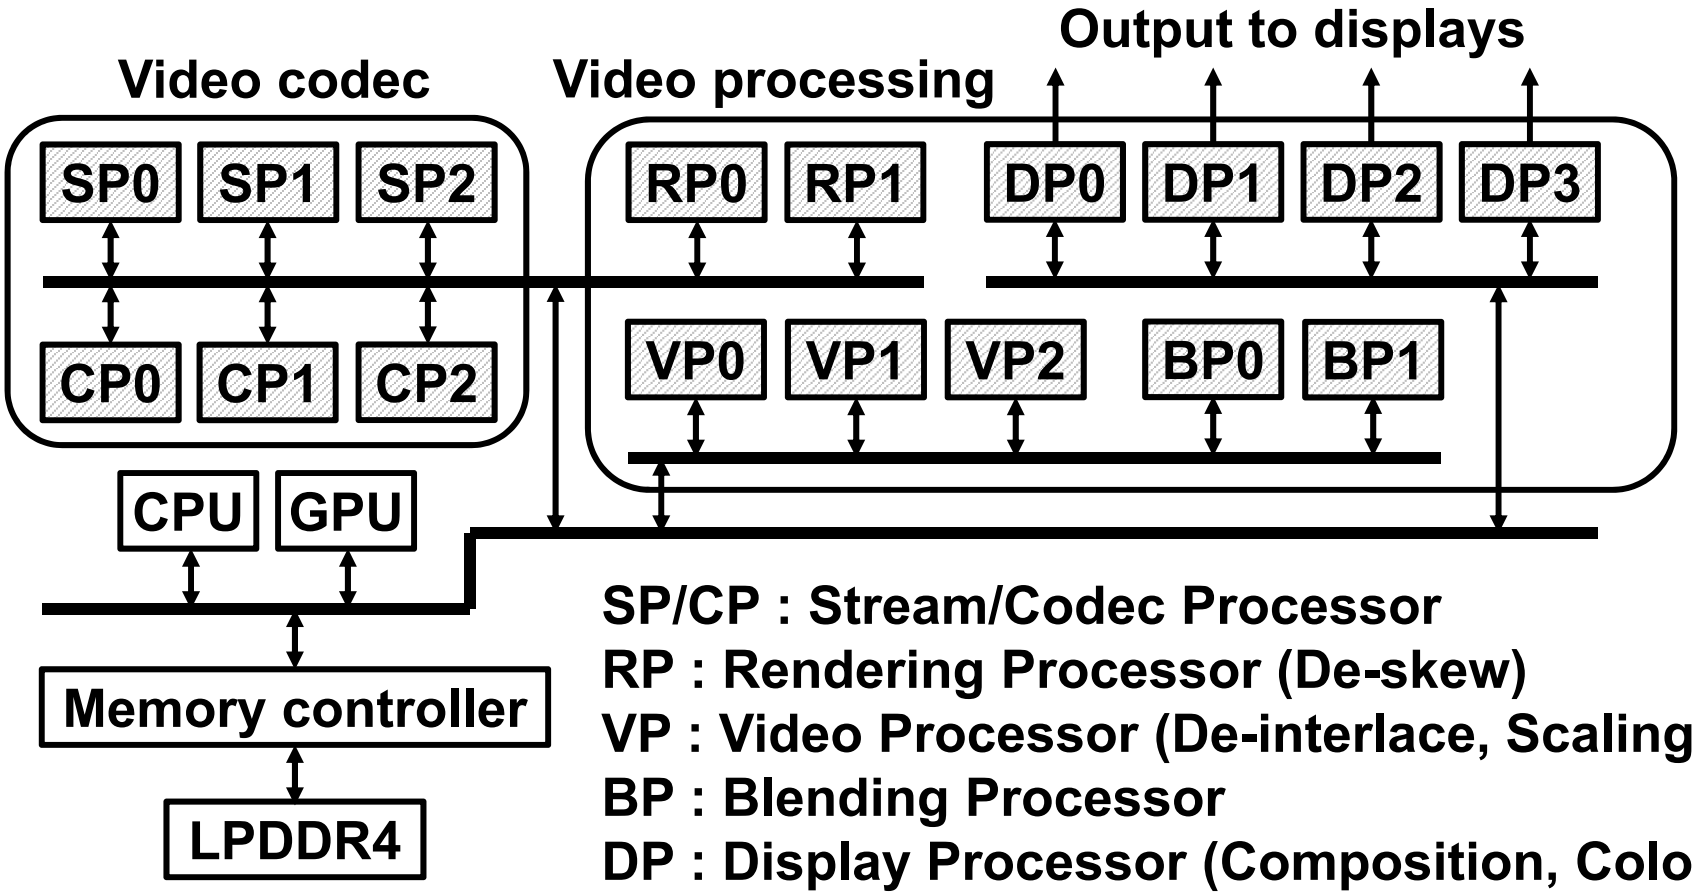

Key technologies in the SoC

- **Full-HD 12-channel video processing at low power, in parallel with cognitive processing**
  - **Specific video processors, operating in parallel with cognitive processing in CPU/GPU.**
  - **Memory bandwidth reduction for video processing**
  - **Peak power control for parallel processing of many video processors.**
  - **A safety mechanism for driver assistance.**

- **Low-latency H.264 decoding and de-skew**
  - **Synchronous and pipelined operation of video processors.**

Key technologies in the SoC

- **Full-HD 12-channel video processing at low power, in parallel with cognitive processing**
  - **Specific video processors, operating in parallel with cognitive processing in CPU/GPU.**
  - **Memory bandwidth reduction for video processing**
  - **Peak power control for parallel processing of many video processors.**
  - **A safety mechanism for driver assistance.**

- **Low-latency H.264 decoding and de-skew**
  - **Synchronous and pipelined operation of video processors.**

Motivation

- **Memory access : Bottleneck of performance and an important factor for low power**
  - **20GB/s in Full-HD 12-channel video processing**
  - **Performance degradation in cognitive processing**
  - **Bus power occupies nearly 40% for video codec.**

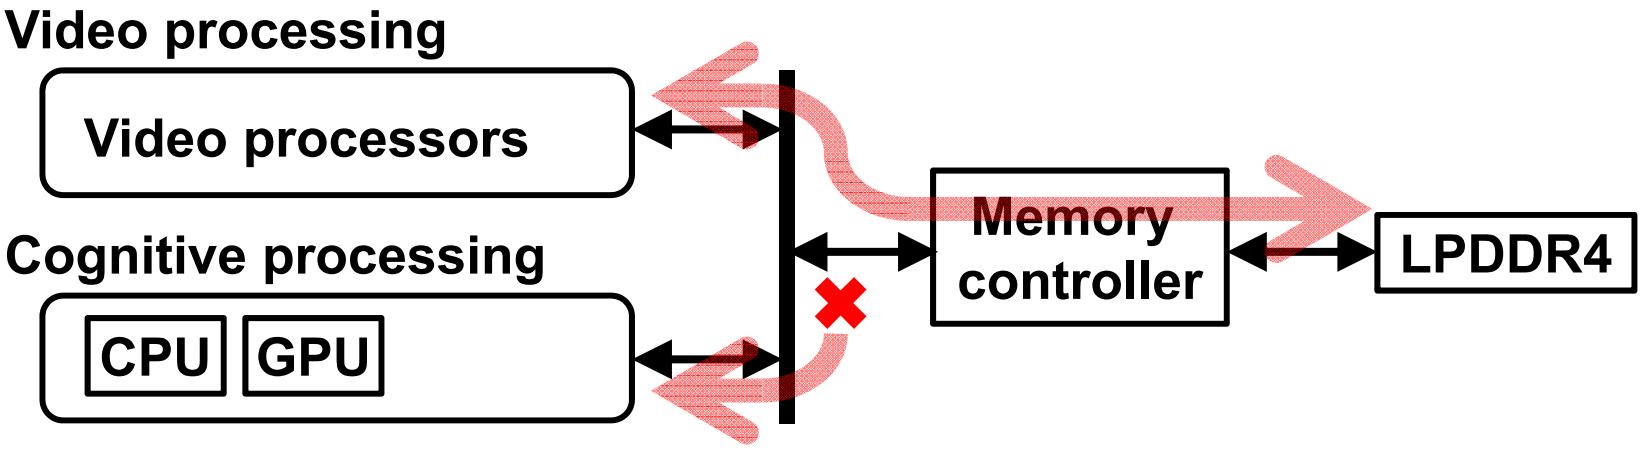

Overview

- Data compression of DDR memory access
  - Reduces memory bandwidth for video processing
  - Helps cognitive processing to access DDR easily
  - Improves total performance of the SoC

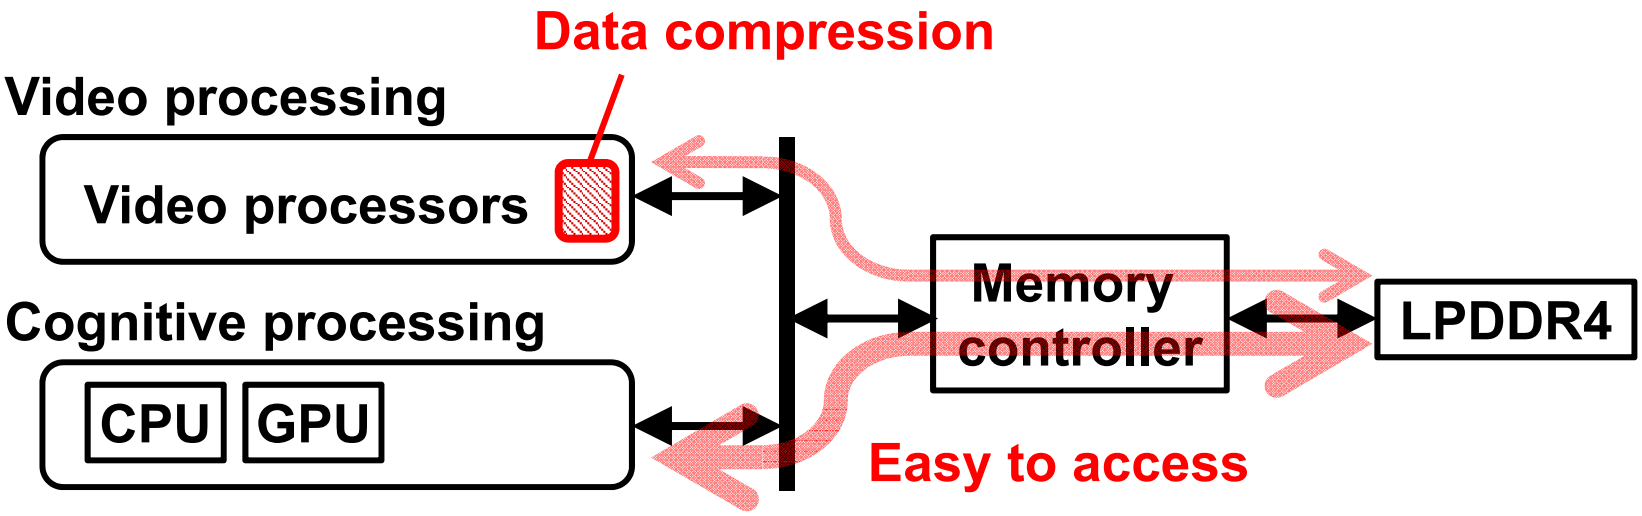

Combination of lossless and lossy

- Balance of image quality and silicon area.
  - Lossless for video codec, lossy for the others.

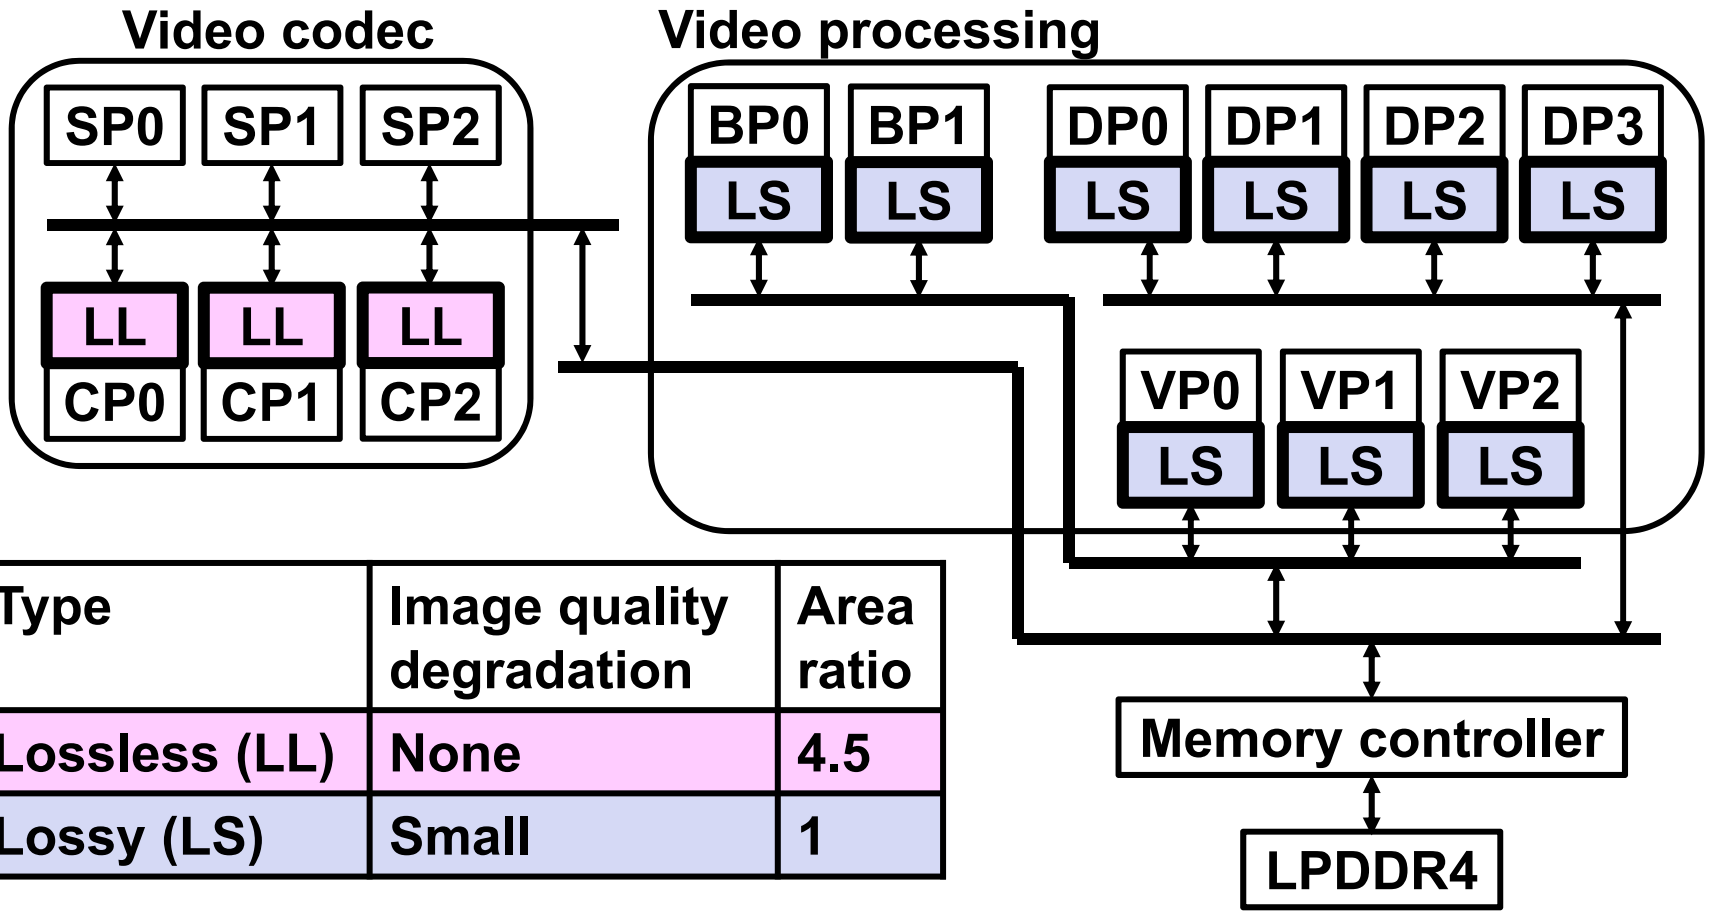

Type	Image quality degradation	Area ratio
Lossless (LL)	None	4.5
Lossy (LS)	Small	1

Lossless and lossy in dataflow

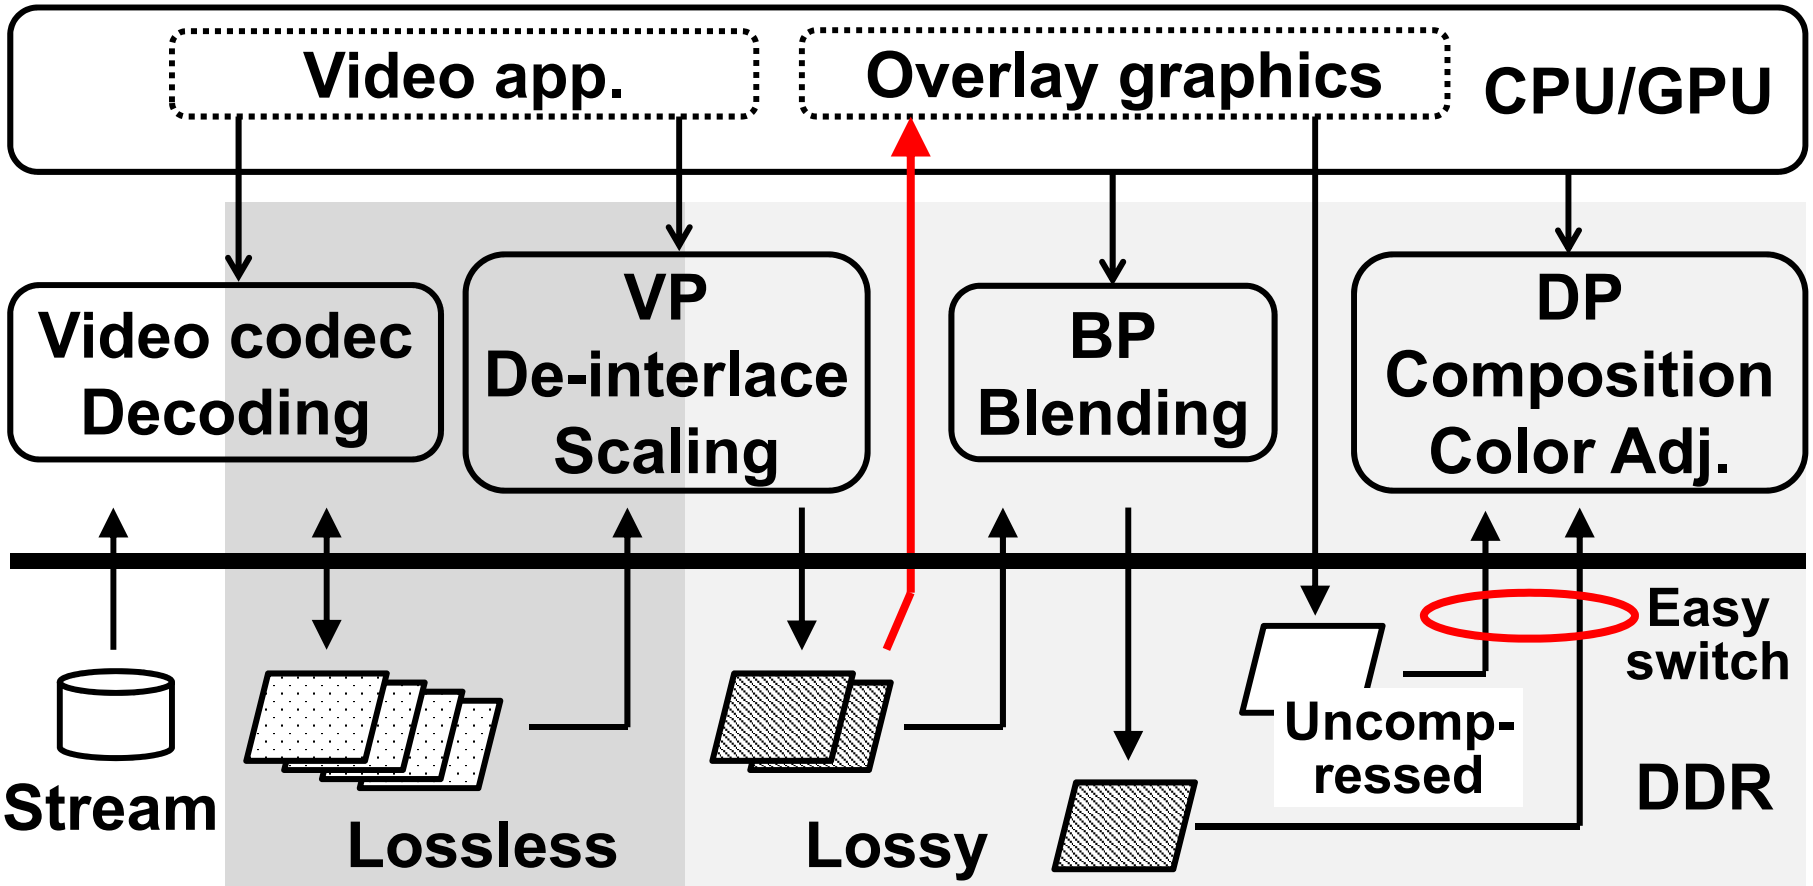

- Issue : CPU/GPU must read lossy-compressed pictures also.

Free access to compressed pictures

- **Lossy-decompression in memory controllers.**
  - Enabled by accessing a specific address space.
  - All modules can access both lossy-compressed and uncompressed pictures with easy switching.

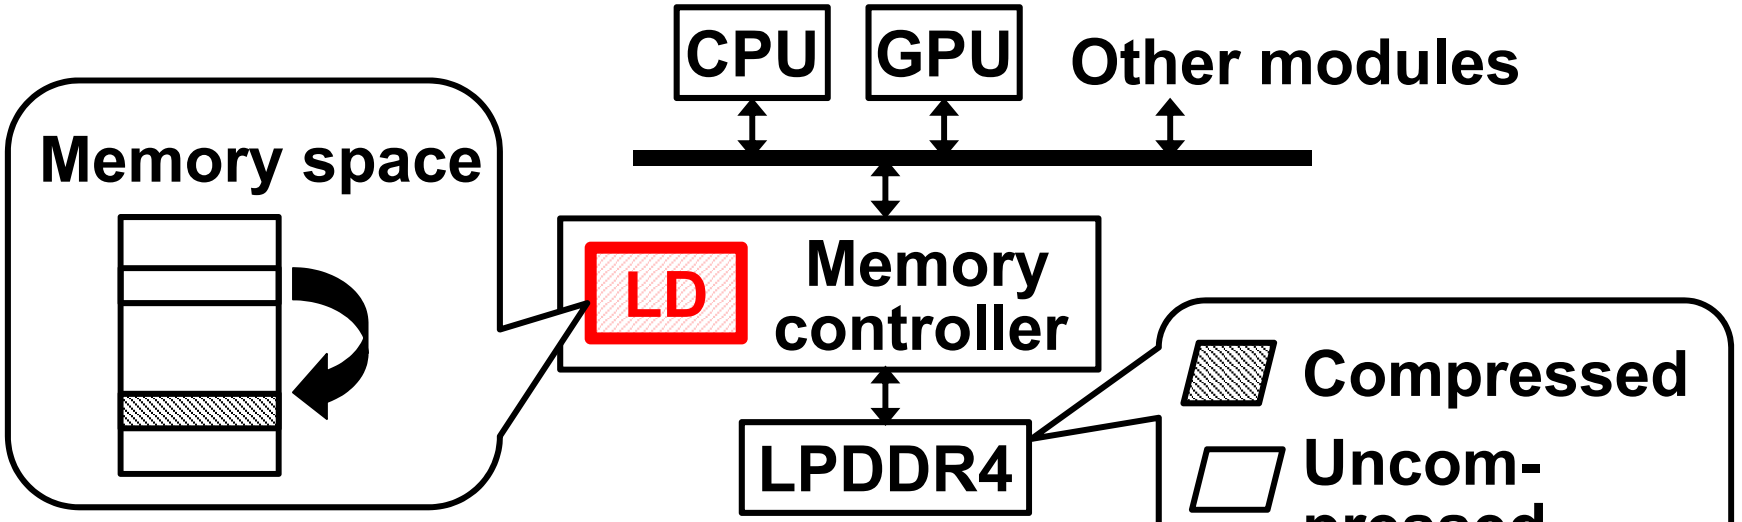

Normal space → LD invalid
Shadow space → LD valid

LD : Lossy Decompression

Data compression methods

- Lossy data compression
- Lossless data compression

Lossy data compression

- Optimized for raster-ordered access
  - Algorithm : DPCM + quantization
- Compression ratio : 50% fixed.
- No visible degradation in typical use-cases.

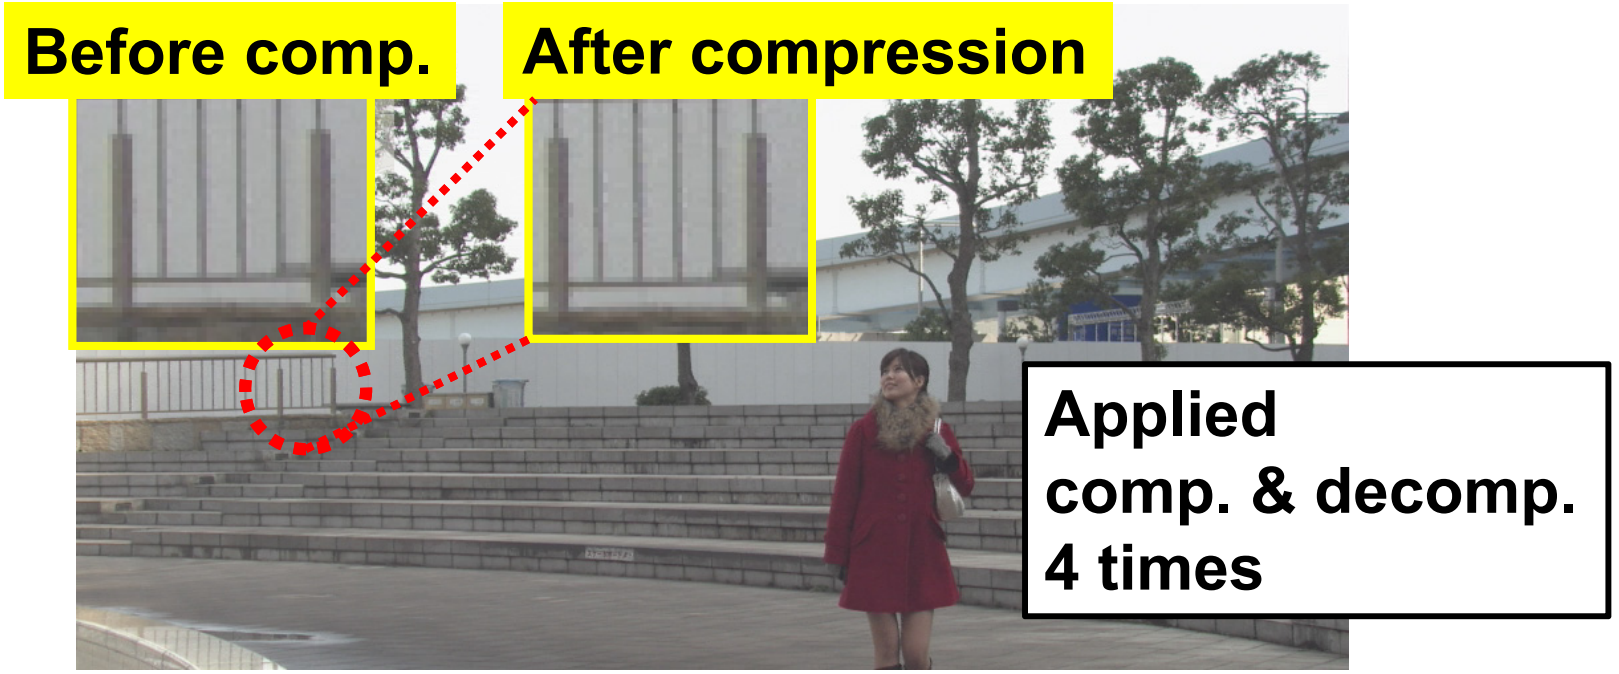

Overview of lossless compression

- **Optimized for a video codec**
  - **Algorithm : Multi-mode 2D-DPCM + VLC**
  - **Picture data is compressed at writing, and decompressed at reading.**

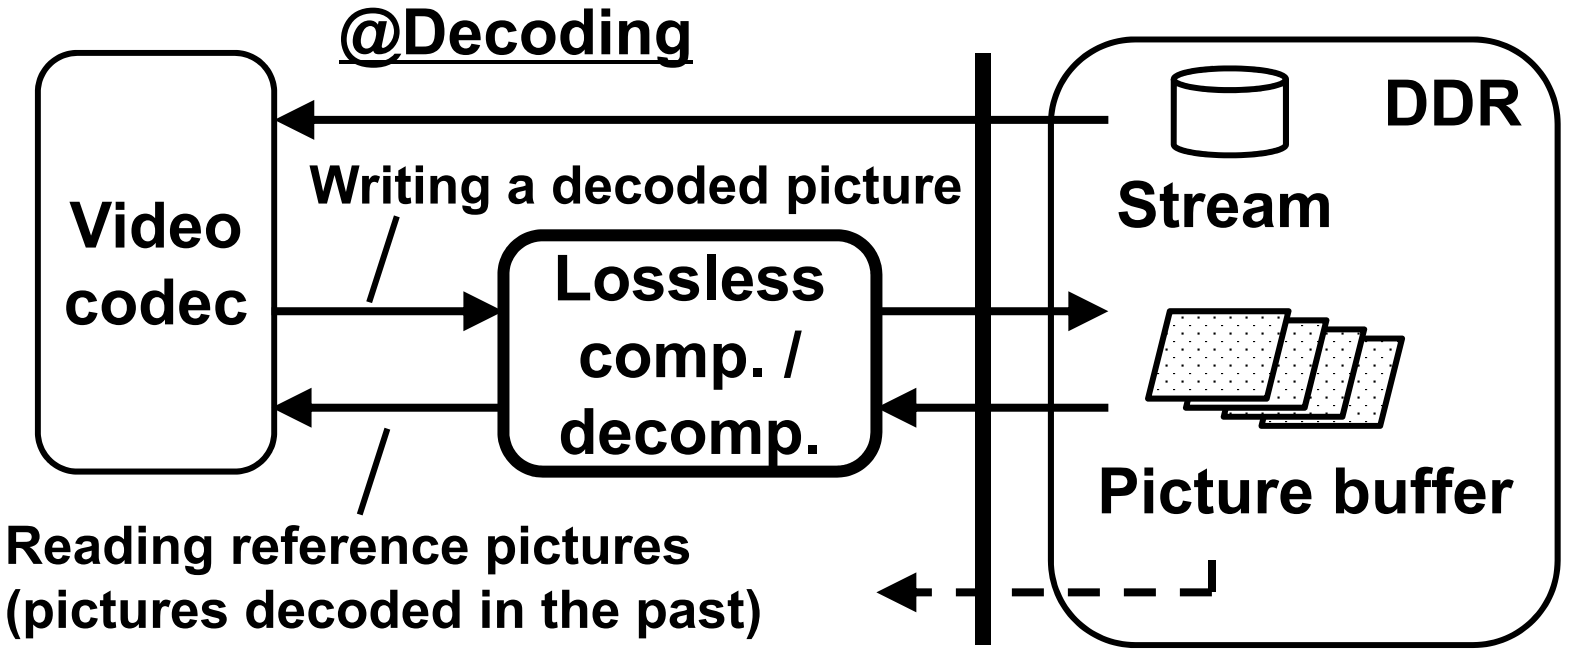

VLC : Variable Length Coding

Many small accesses from the codec

- **Picture R/W from the video codec**
  - **In unit of divided square areas of a picture.**
    - **Small (about 10B-20B) access every line**

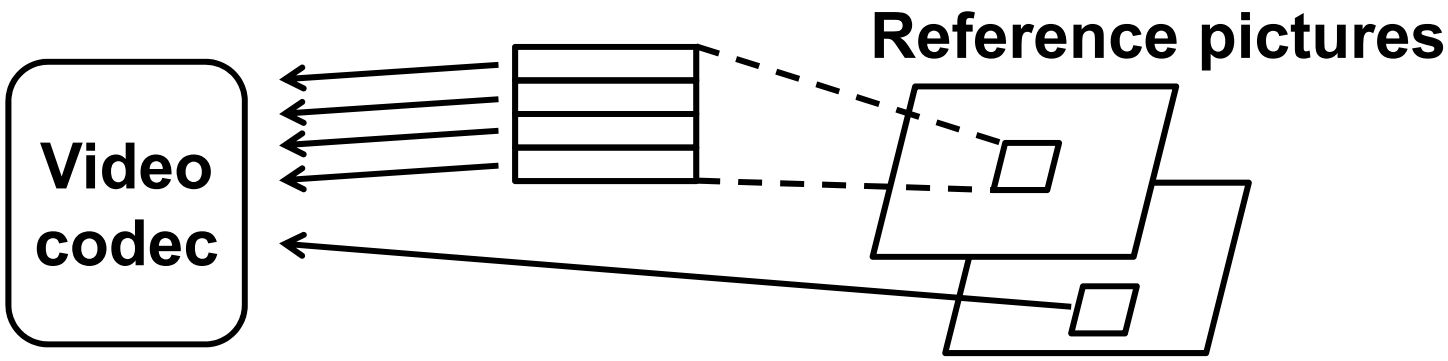

- **Conventional method**
  - **Such small data is simply compressed to smaller data for memory bandwidth reduction.**
  - **A severe problem as follows.**

Dilemma in data compression

- **Smaller data, lower DDR-access efficiency**
  - Consecutive accesses of small data size
    - Overheads due to the timing restriction between the DDR commands.

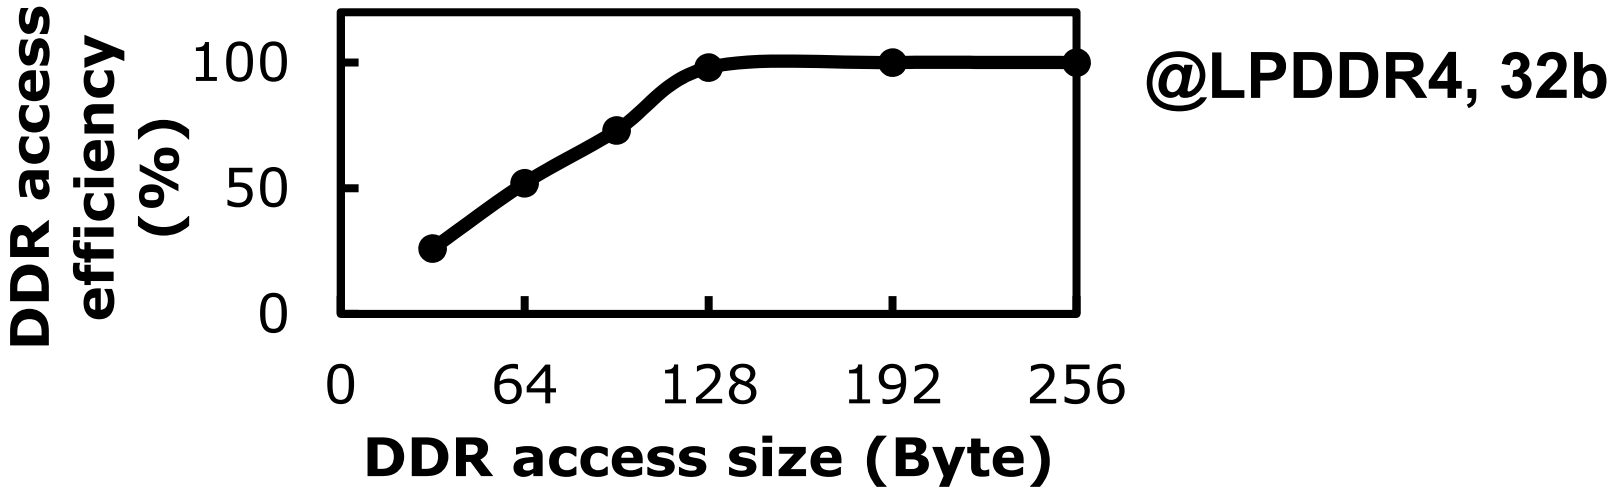

Dilemma in data compression

- **Ex. 50% fixed data compression**
  - **Case A : 256B to 128B**
    - **50% memory bandwidth is achieved.**

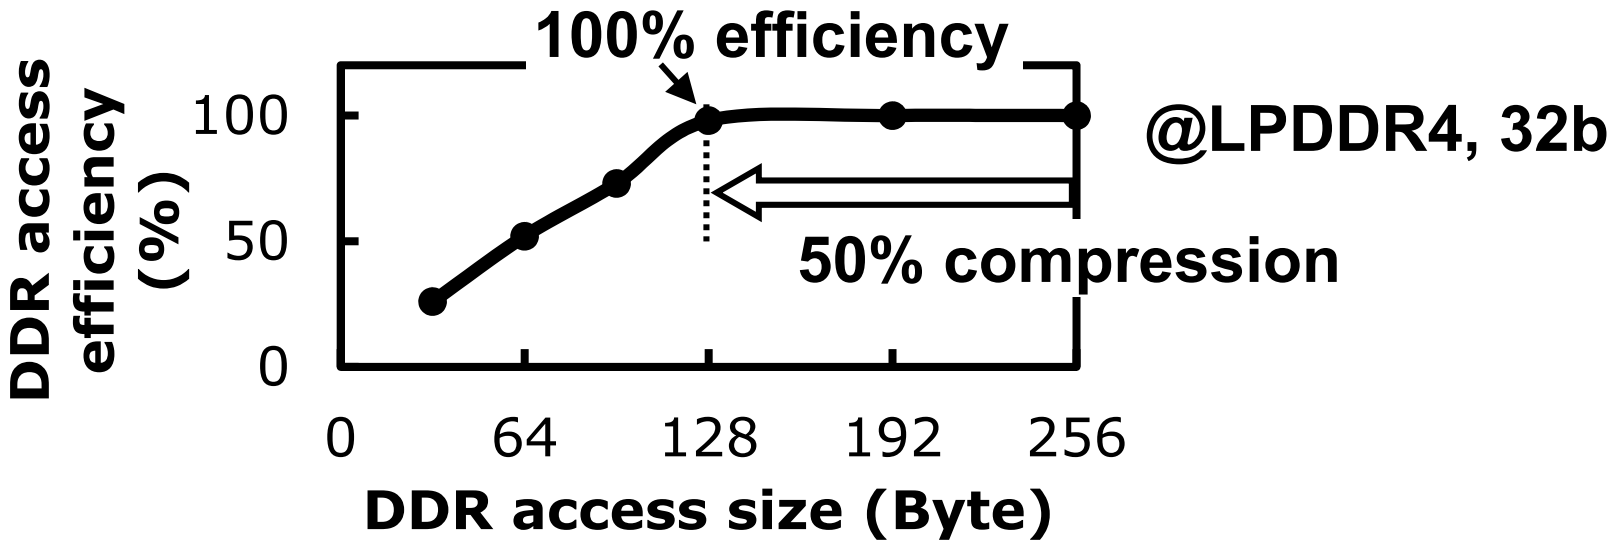

Dilemma in data compression

- **Ex. 50% fixed data compression**
  - **Case B : 128B to 64B**
    - **The consumed memory bandwidth is not changed because of 50% access efficiency at 64B.**

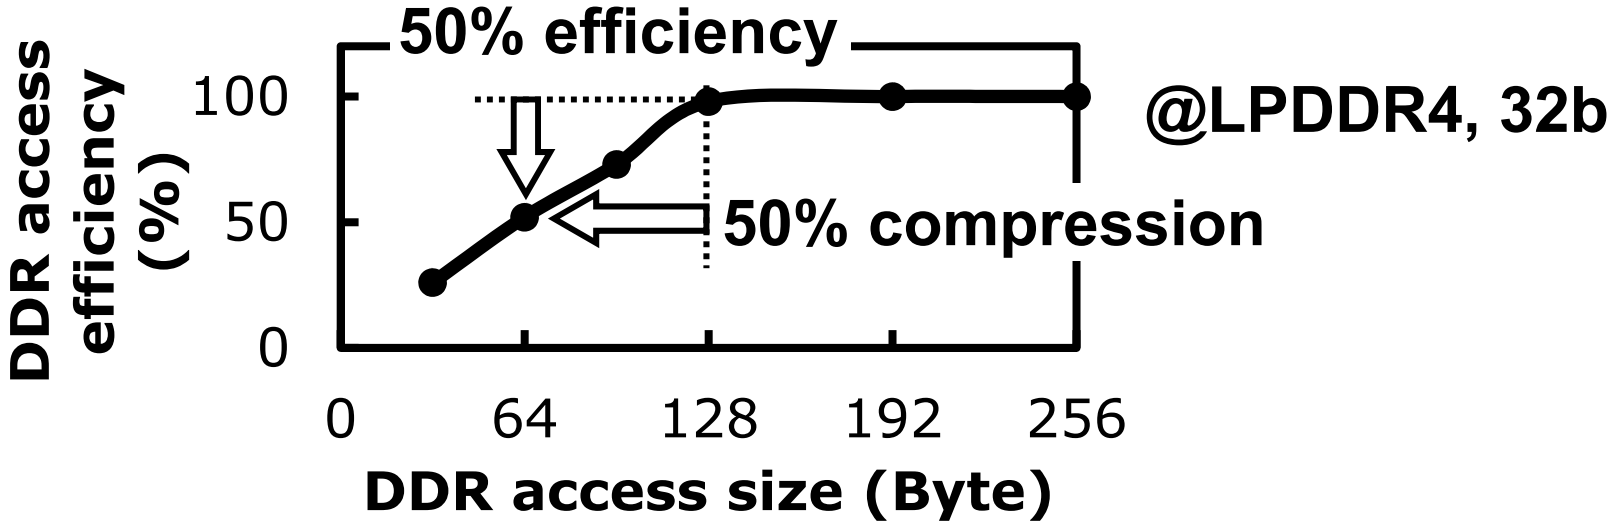

Dilemma in data compression

- **Data compression in small unit**
  1. Small data becomes smaller after compression.
  2. DDR access efficiency becomes lower.
  3. Memory bandwidth is not reduced enough.

➔ **Enough compression effect is not expected in the conventional method.**

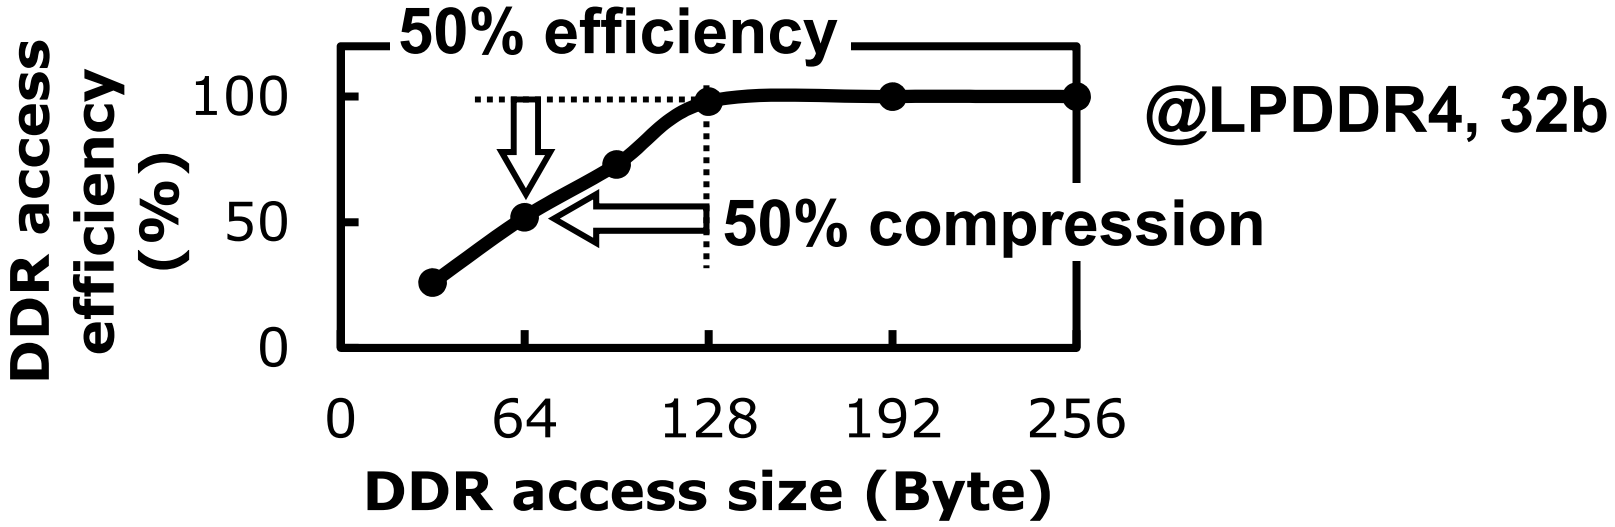

Proposed architecture

- **A write buffer and a cache** in the compression module to enlarge memory access size.
  - Execution unit of the compression module : 256B
  - The write buffer unifies small data to 256B.
  - Compressed to multiples of 64B : 256/192/128/64B
  - Decompressed 256B is stored to the cache.

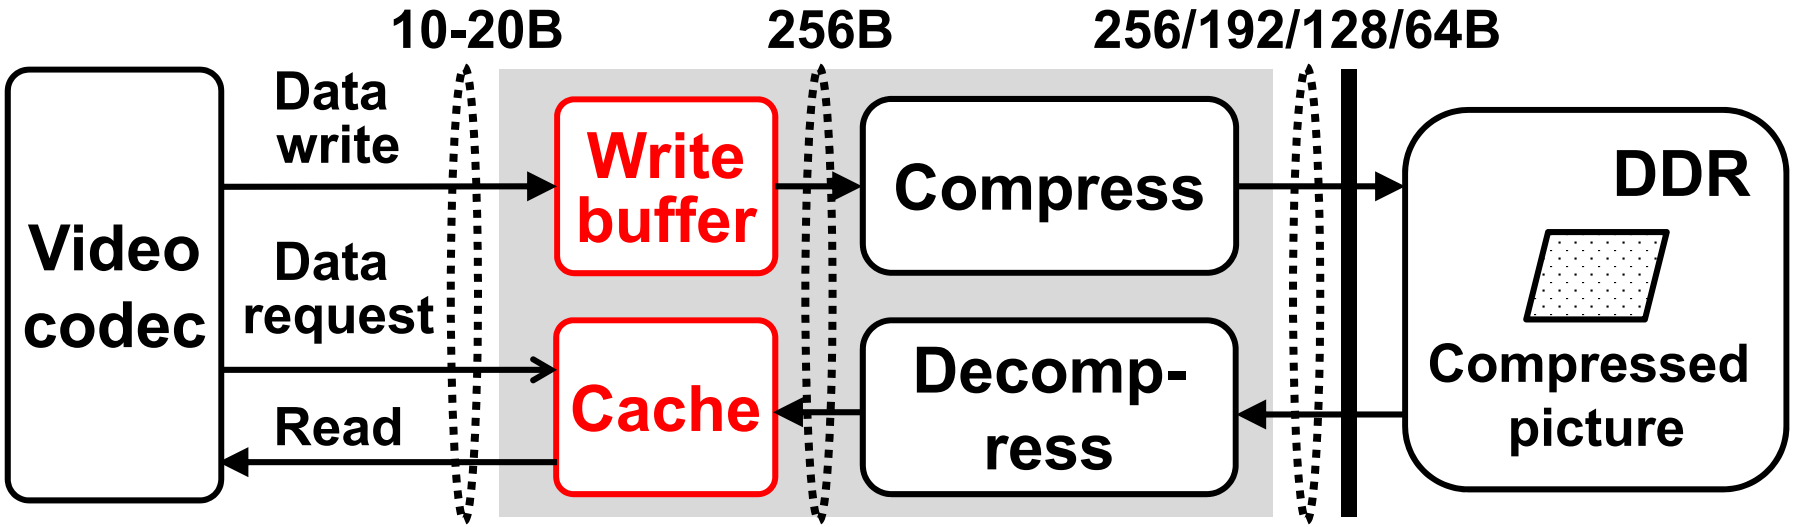

Select execution unit by simulation

- **Compression ratio =**

$$\frac{\textit{Traffic from/to DDR with considering efficiency}}{\textit{Traffic from/to the video codec}}$$

  - **Smaller is better. 100% without compression.**

256B is better than smaller unit size

- **Reason of better result at large execution unit**
  - **The compression ratio is affected by low DDR access efficiency at the size after compression.**

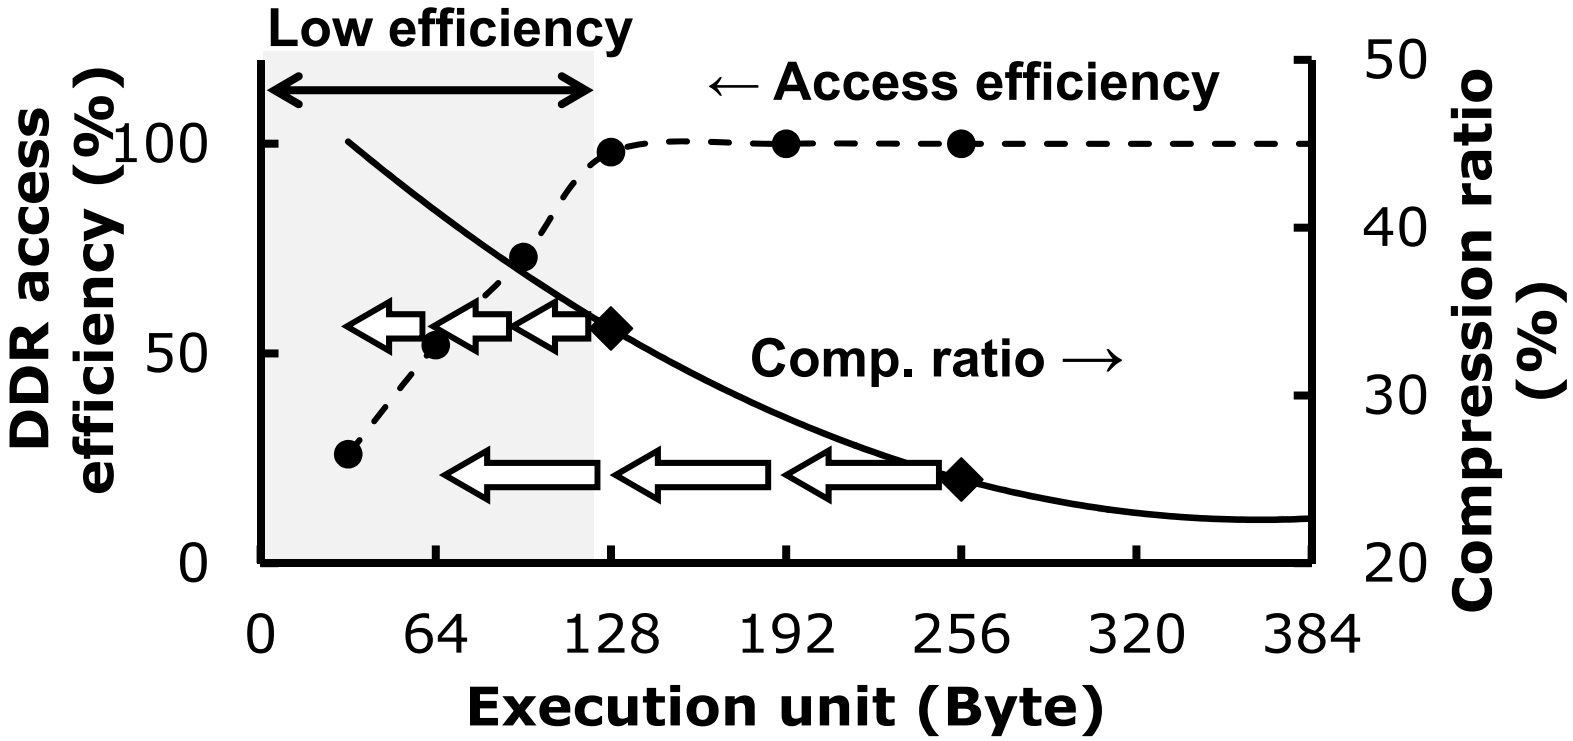

256B is the best unit size for LPDDR4

- How about much larger size?
  - Large entry size of the cache causes less number of entries under the constant cache-memory size.

  ➔ More cache-miss overhead.

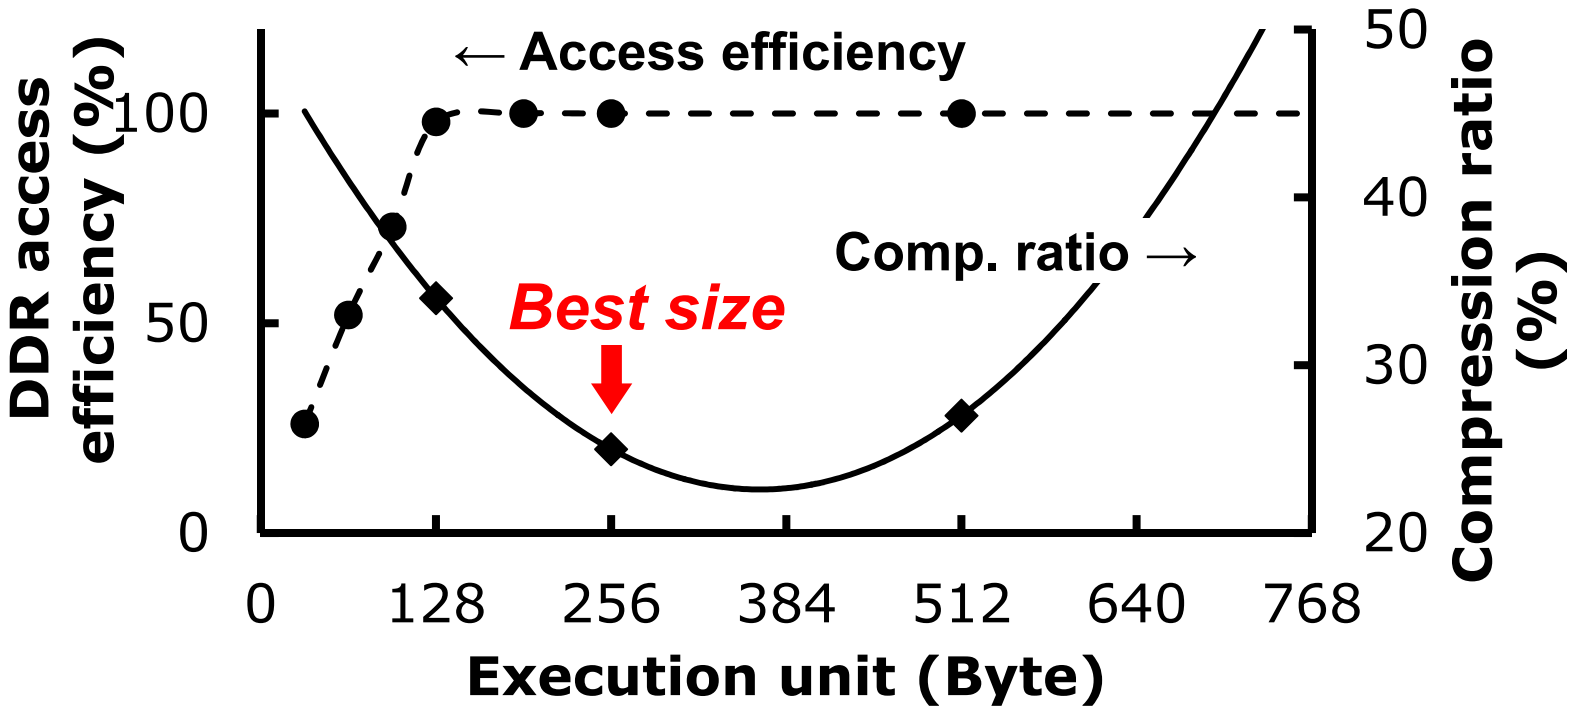

Experimental results (lossless)

- Compression ratio of lossless compression
  - 70% reduction of memory bandwidth for picture read/write in video codec processing.
  - Cf. Without a write buffer and a cache : 50% or less reduction

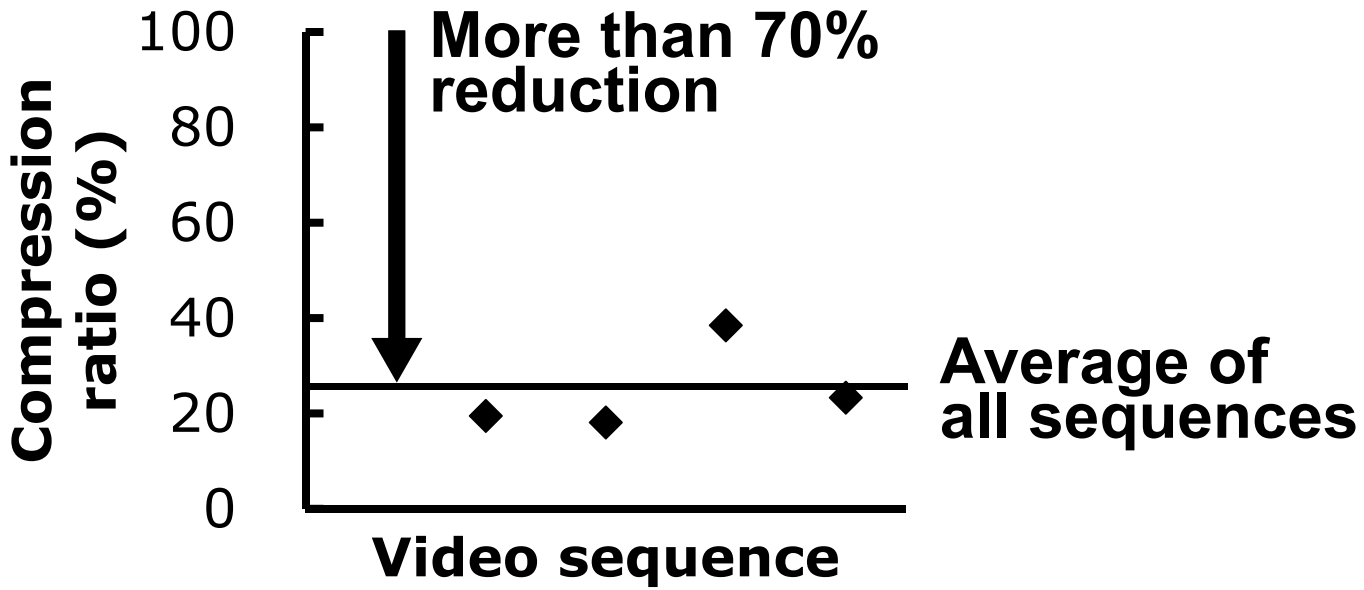

Experimental results (lossless+lossy)

- **Memory bandwidth in typical Full-HD 12-channel processing**
  - **Lossless for video codec, lossy for the others.**
  - **50% reduction of memory bandwidth in total.**

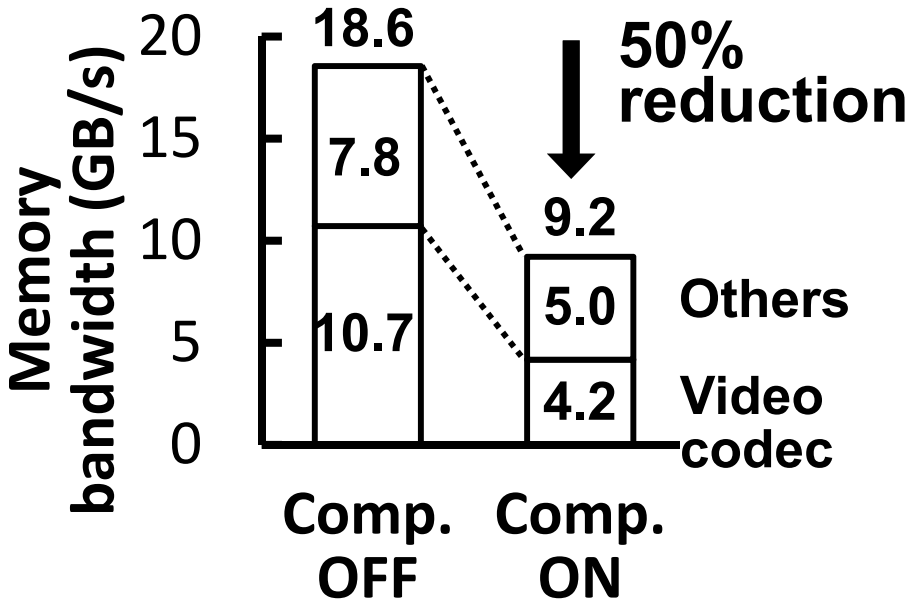

Experimental results (power)

- Power consumption in Full-HD 12-channel H.264 decoding : 197mW
- 20% reduction by data compression
  - Corresponds to a half of bus power.

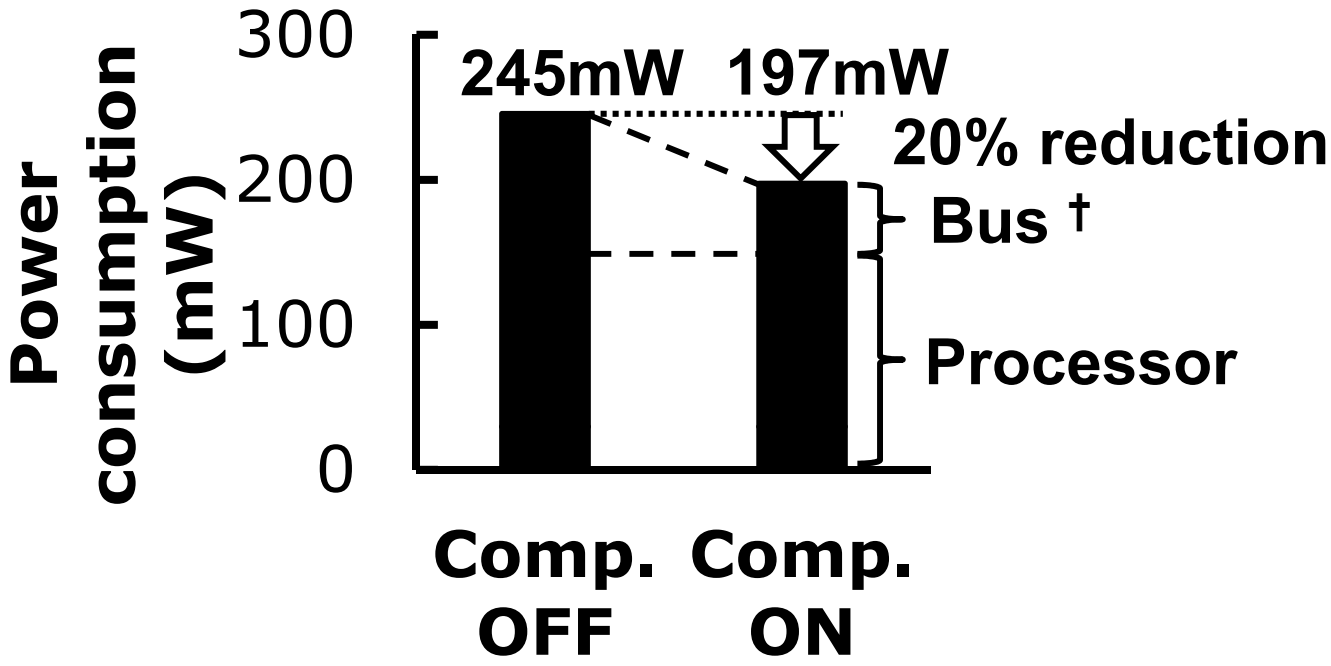

† Estimated from the ratio of power consumption for I-PIC and P-PIC

Comparison with previous works

- Power efficiency of decoding H.264 and HEVC streams : 0.16-0.29nJ/pixel (avg. 0.24nJ/pixel)

	This work	ISSCC 2015	ISSCC 2013	ISSCC 2012	ISSCC 2009
Chip type	AP	AP	Decoder	AP	AP
Technology	16n 0.8V	28n 0.9V	40n 0.9V	32n 1.1V	65n 1.2V
Performance	750 Mpix/s	249 Mpix/s	249 Mpix/s	124 Mpix/s	62 Mpix/s
Core power	0.16-0.29 nJ/pix	0.51 nJ/pix	0.31 nJ/pix	0.52 nJ/pix	5.50 nJ/pix

AP : Application processor

Conclusion

- The SoC for integrated cockpit systems.
  - Implemented in 16nm FinFET.

- Full-HD 12-channel video processing by 17 video processors of 6 different types.
- 50% memory bandwidth reduction by lossless and lossy data compression.
- 70ms-latency H.264 decoding and de-skew.

- 197mW Full-HD 12-channel H.264 decoding.
  - 20% power-saved by memory bandwidth reduction.

Demonstration at DS1

- From 5 pm until 7 pm TODAY.
- In the Golden Gate Hall.

- Proof of Full-HD 12-channel H.264 playback
  - Lossless memory-bandwidth reduction : -70%

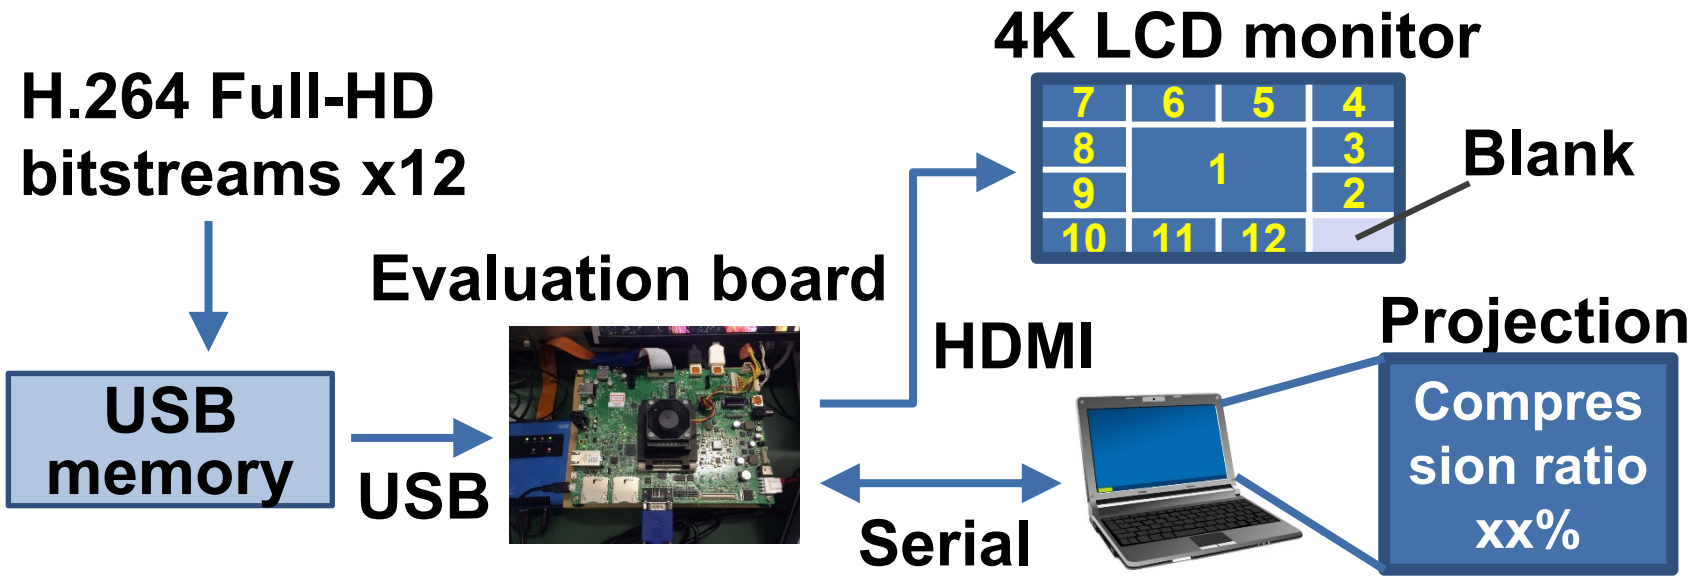

Thank you.

A 16nm FinFET Heterogeneous Nona-Core SoC Complying with ISO26262 ASIL-B: Achieving 10^{-7} Random Hardware Failures per Hour Reliability

C. Takahashi[1], S. Shibahara[1], K. Fukuoka[1],
J. Matsushima[1], Y. Kitaji[1],
Y. Shimazaki[2], H. Hara[2], T. Irita[2]

[1] Renesas System Design
[2] Renesas Electronics

Outline

- ISO26262
- High Flexibility Run-time Self Test
- Killer-droop Monitor with Droop Prediction
- Conclusion

ISO26262
Road vehicles - Functional safety

- An adaptation of the Functional Safety standard IEC 61508 for Automotive Electric/Electronic Systems
- Automotive-specific risk-based approach
- Defines “Automotive Safety Integrity Levels (ASIL)”

e.g. Failure Rate Target

Safety Level

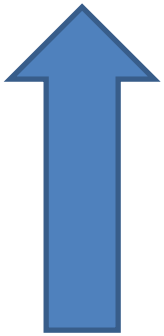

ASIL	Target rate
D	< 10^{-8} per hour (≈ 1 per 11408 years)
C / B	< 10^{-7} per hour (≈ 1 per 1141 years)
A	< 10^{-6} per hour (≈ 1 per 114 years)

Strategy against Random Hardware Fault

Safety mechanism detects random hardware faults

→ Go into safe state within
"Fault Tolerant Time Interval" (FTTI)

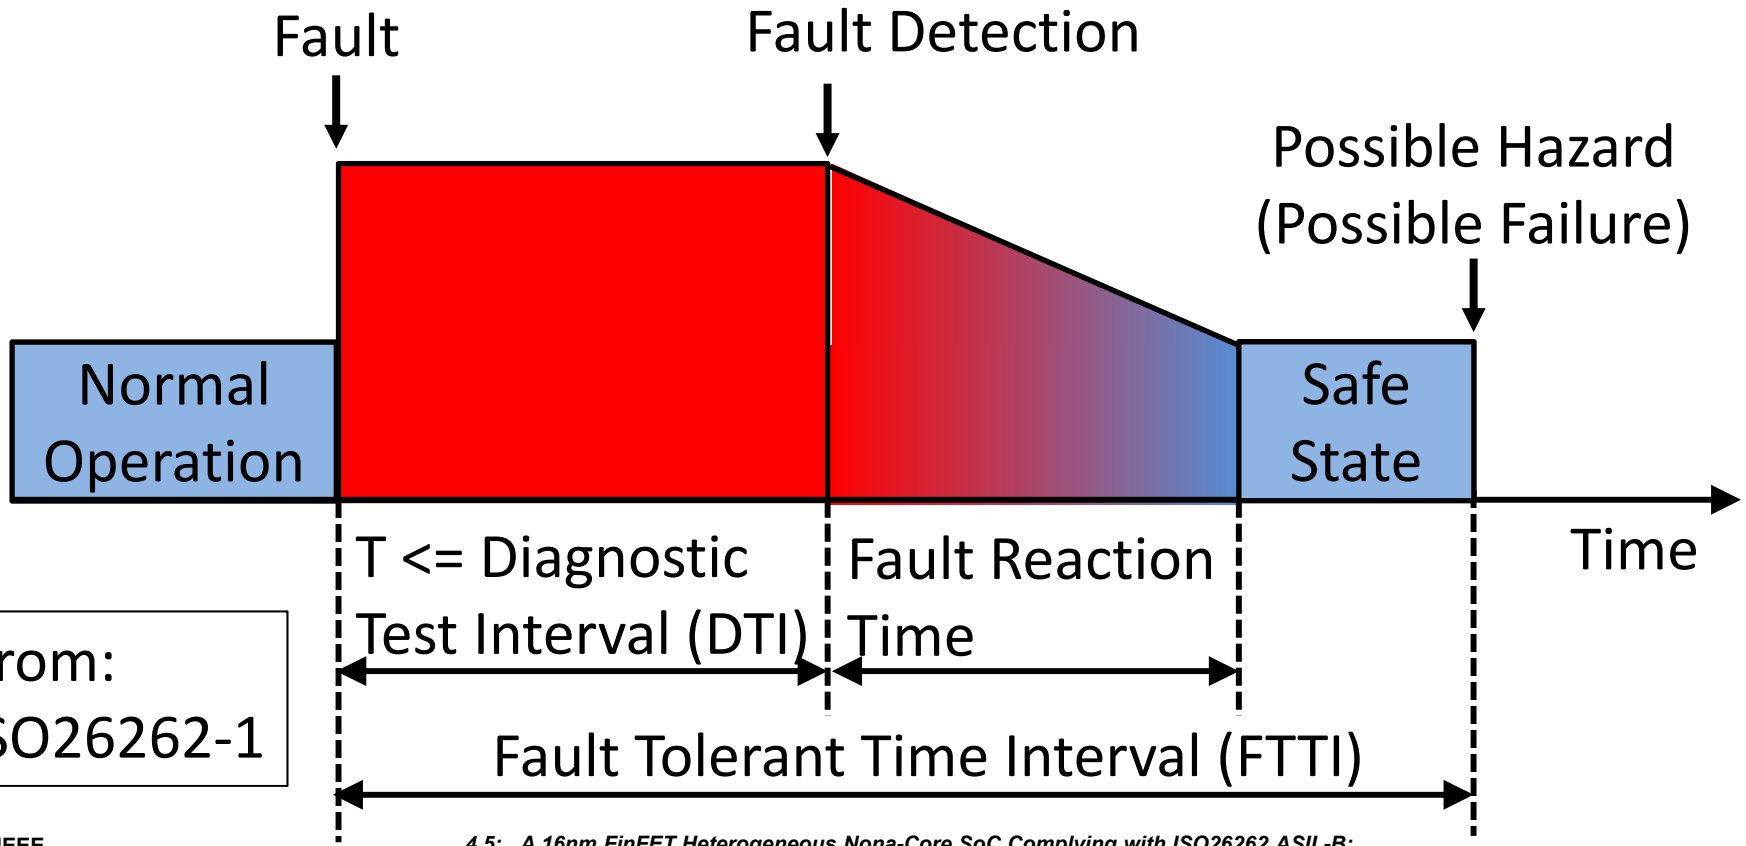

From:
ISO26262-1

What is the matter in Safety Mechanism?

- Extra cost for Safety Mechanism
  - Software execution time (e.g. Software self test)
  - Logic size (e.g. Redundant logic)
  - Power consumption

Growing chip complexity increases the cost!

- Discontinuous time of the original functionality
  - Stops running program/calculation
  - Stops CPU interrupt response

Automotive is moving even in self test!

- Detection oriented method for random H/W fault.
  - ISO26262 assumes “detection” and “reaction”.

“reaction” cannot avoid “reaction” cost!

Our Solution

1. High flexibility run-time self test
   - Utilizes Built-In Self Test (BIST) hardware to detect random hardware fault
   - Supports time slicing run-time self test
   - Supports individual run-time self test

Provides fault detection mechanism with minimum cost

2. "Killer-droop" monitor with droop prediction
   - High speed voltage monitoring
   - Voltage prediction unit
   - 0-cycle delay clock control

Provides fault prediction for voltage droop

Die Photo and Specifications

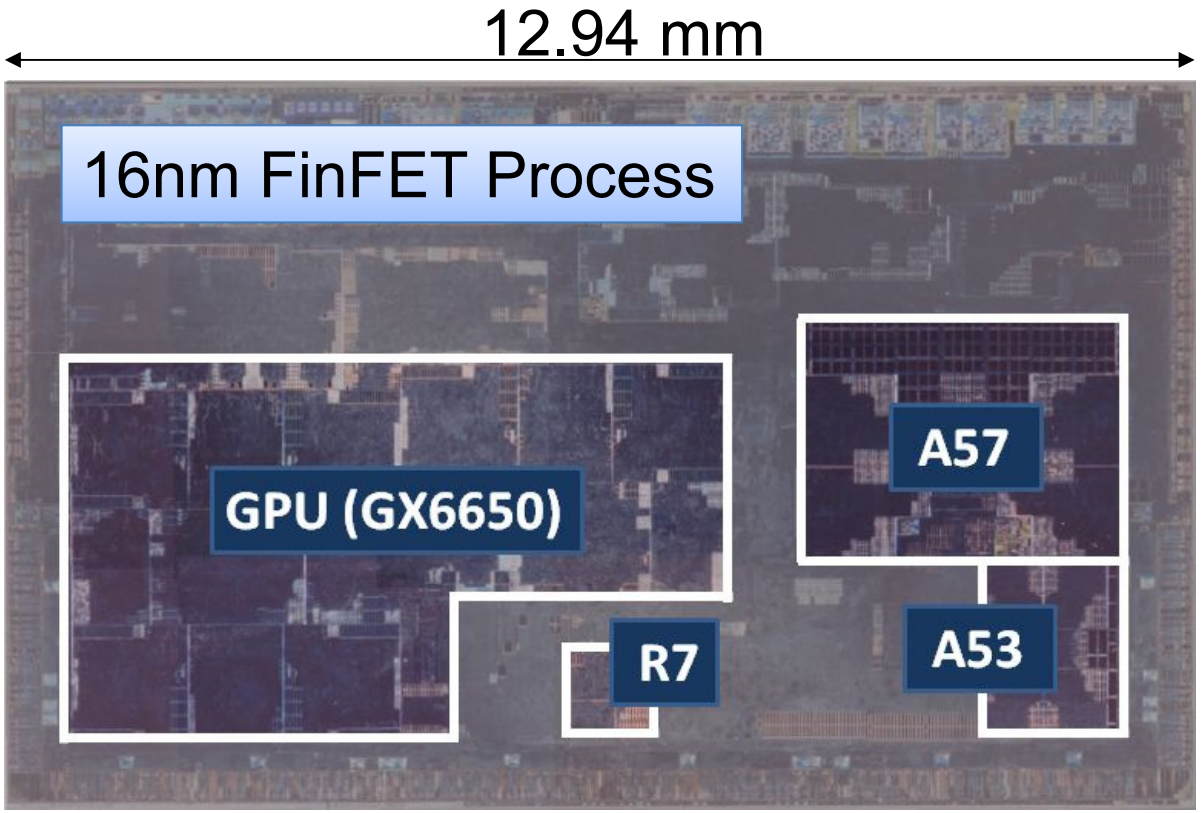

CPU	ARM Cortex-A57 Quad	2.0 GHz	(1)(2)
	ARM Cortex-A53 Quad	1.2 GHz	(1)
	ARM Cortex-R7 Dual-core Lock-step	800 MHz	
GPU	Imagination Technologies GX6650	700 MHz	(1)

High Flexibility Run-time Self Test

- Implements BIST logic in each functional block
- Implements control logic for:
  - Time slicing run-time self test
  - Individual run-time self test

Satisfies the ISO26262 ASIL B Diagnostic Coverage within Diagnostic Test Interval (DTI)

Block Diagram

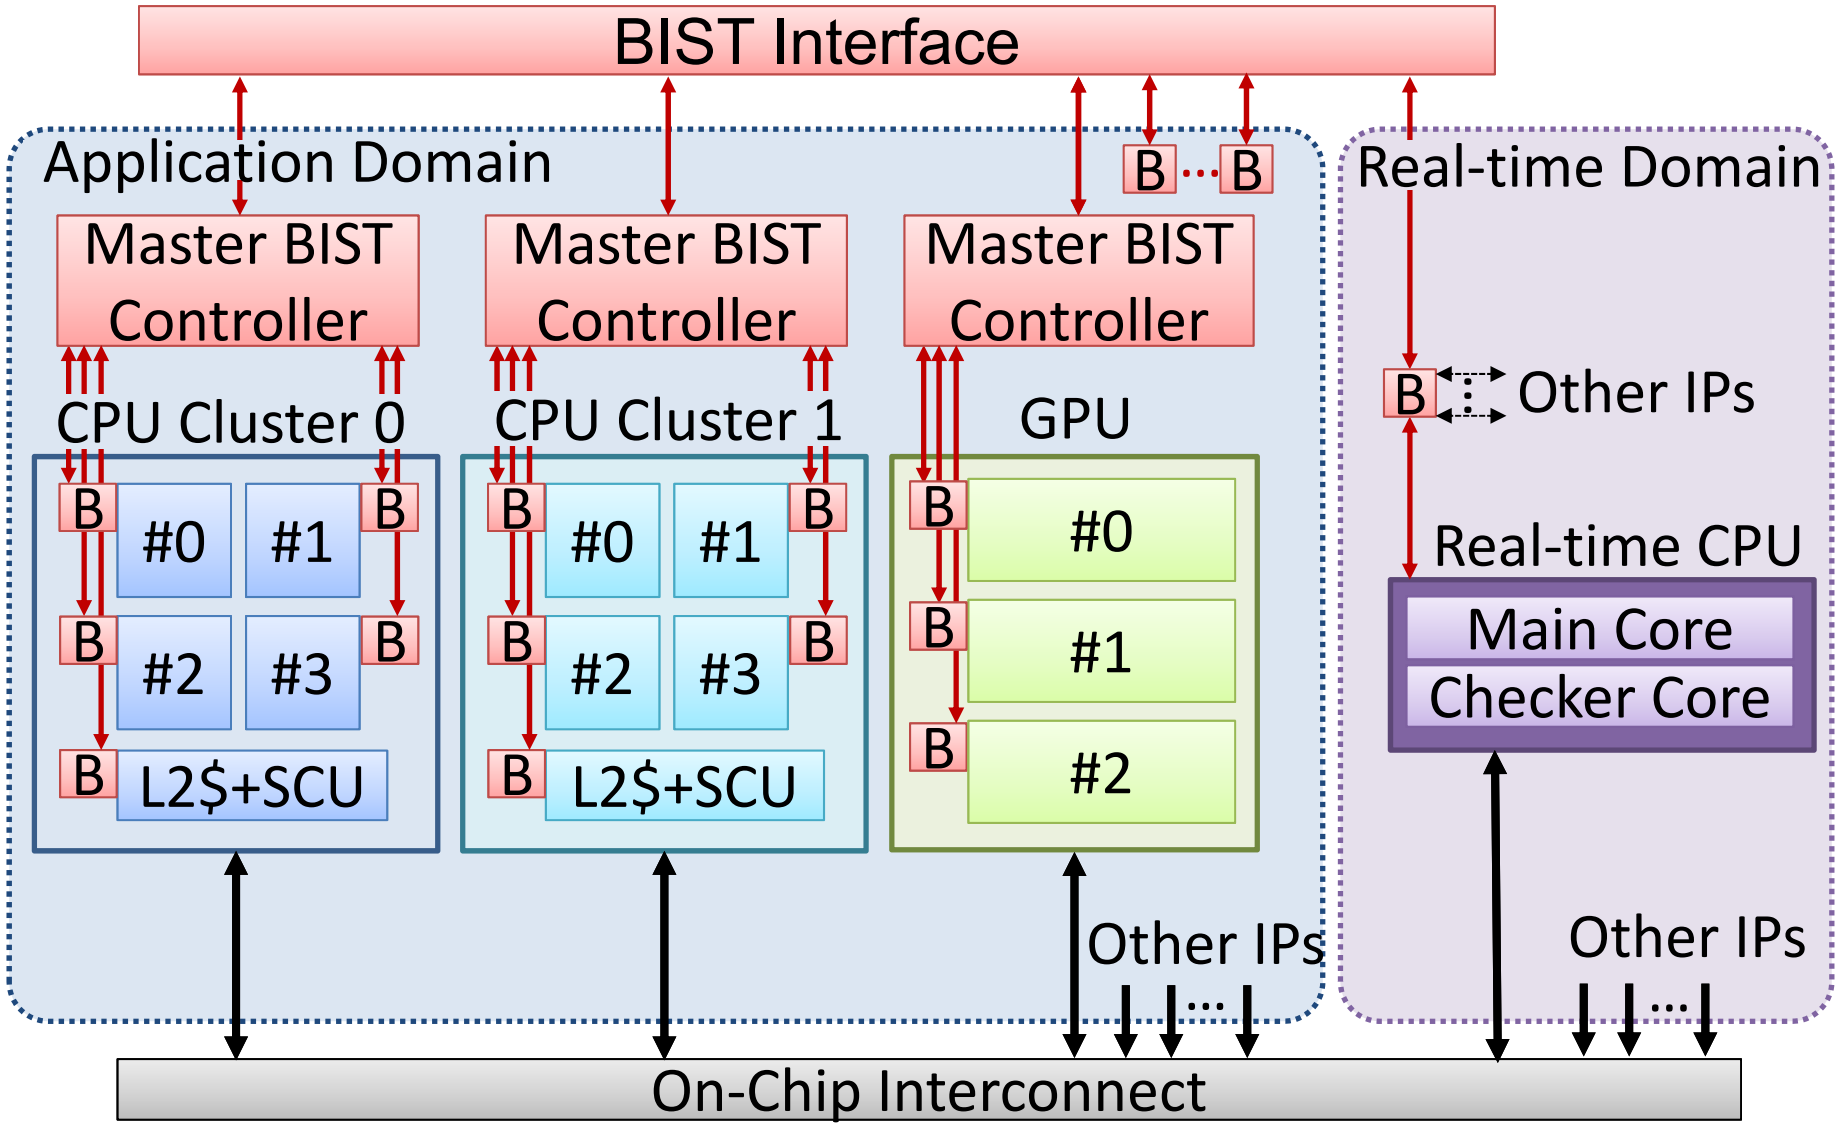

Implementation for CPU Cluster

- Cores waits for L2$+SCU self-test while in sleep state

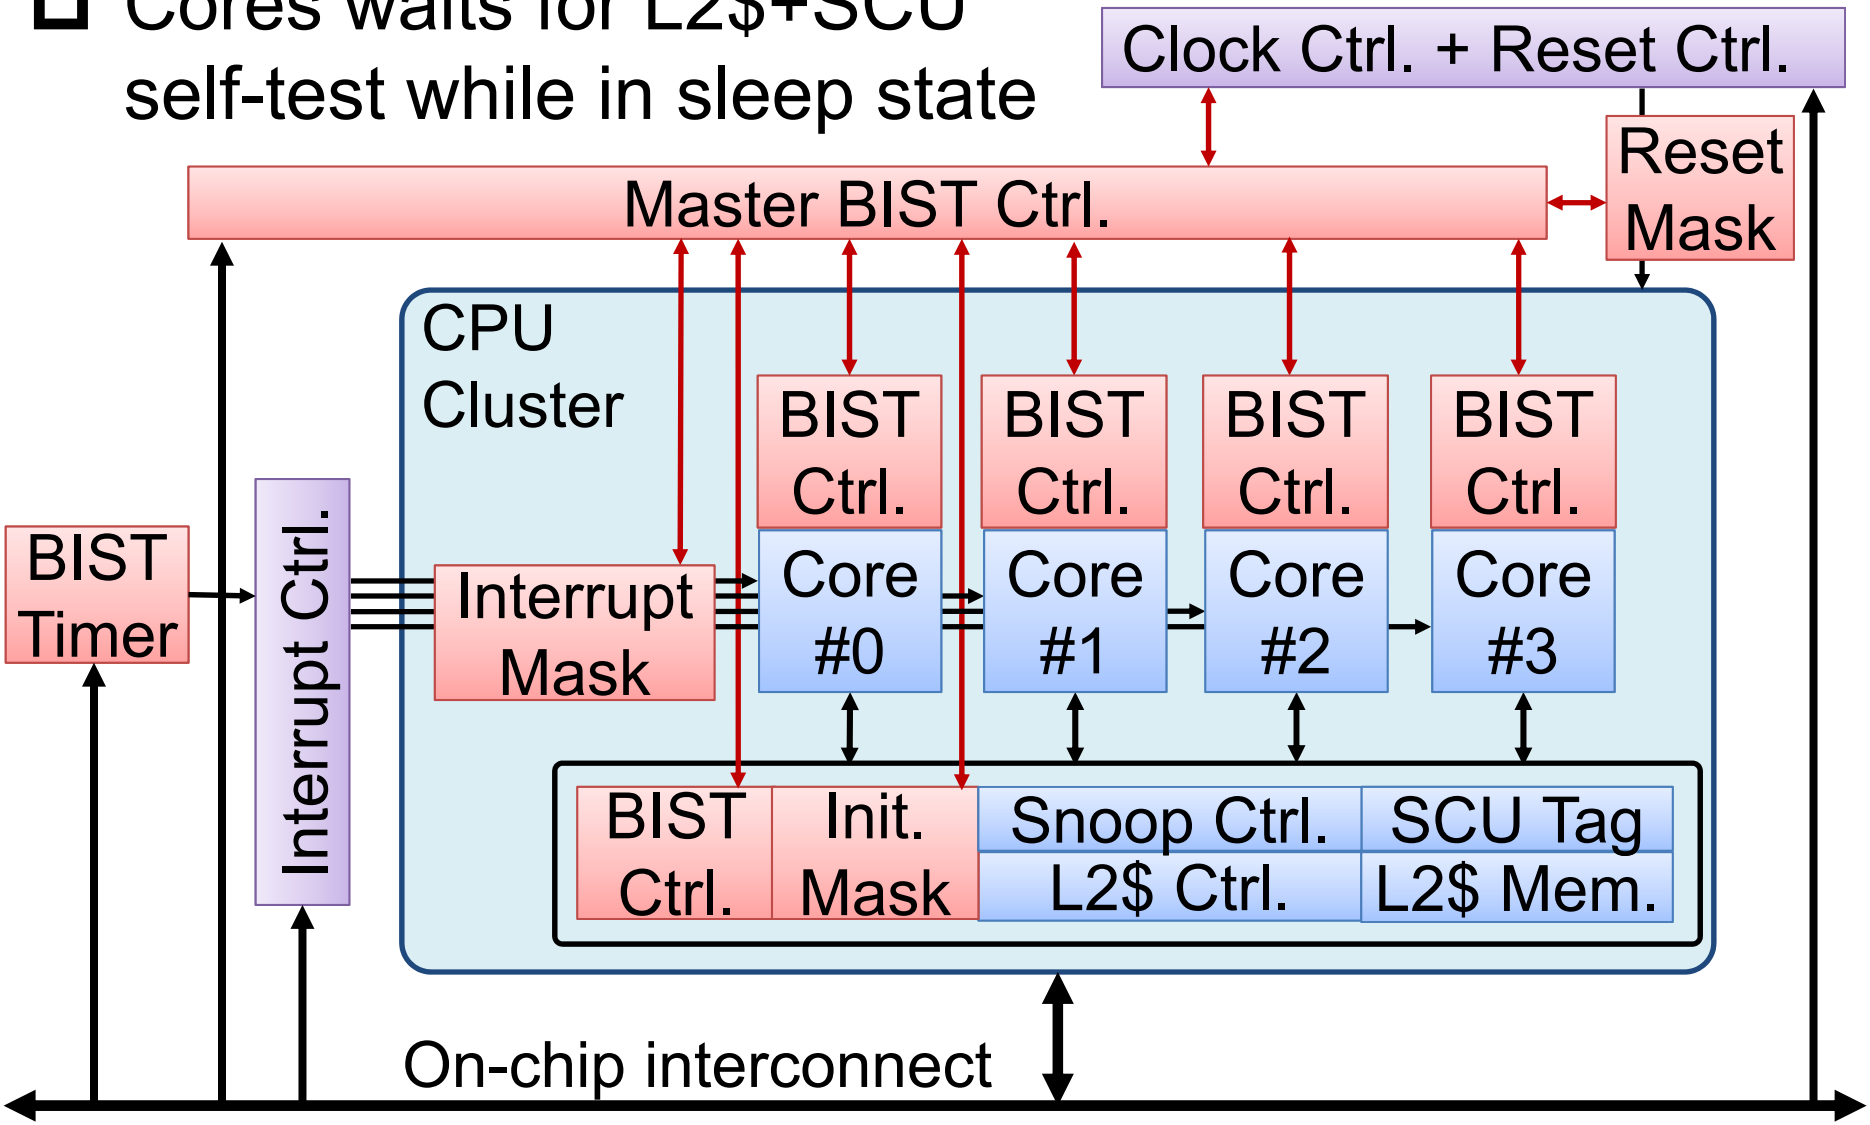

Slicing Run-time Self Test

- Time Slicing
- Module independent test

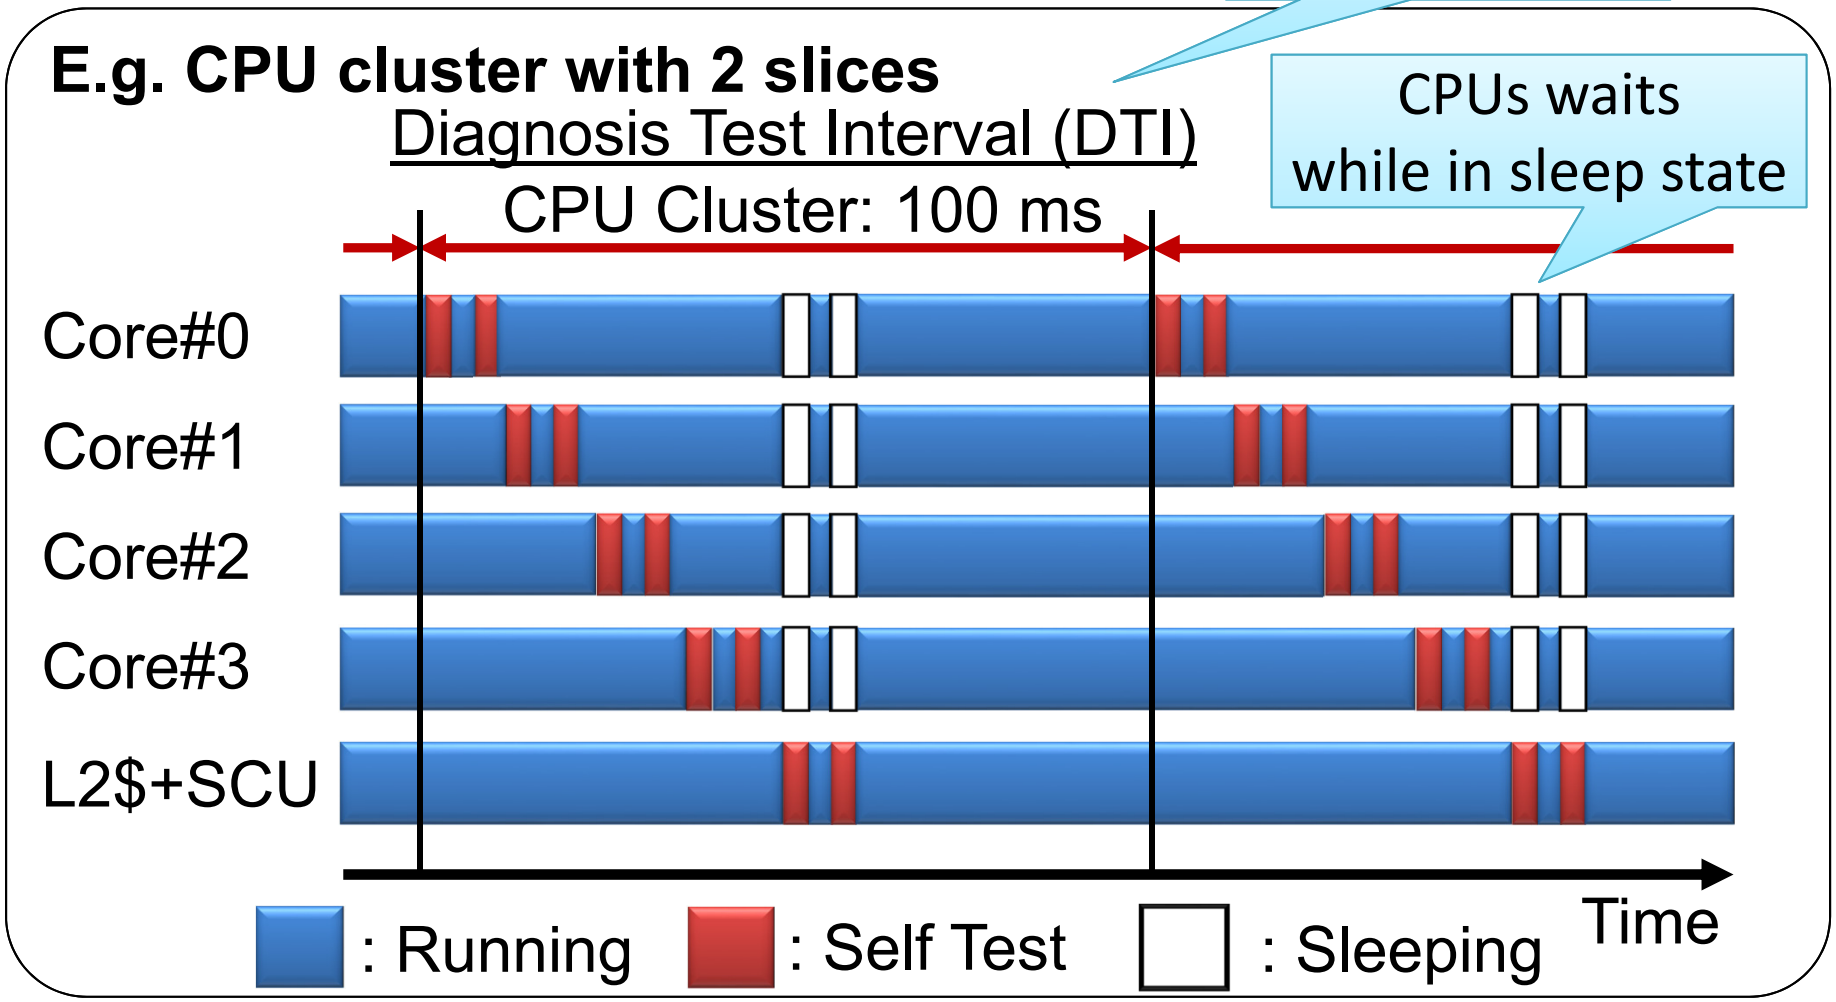

Test Time Slicing Result

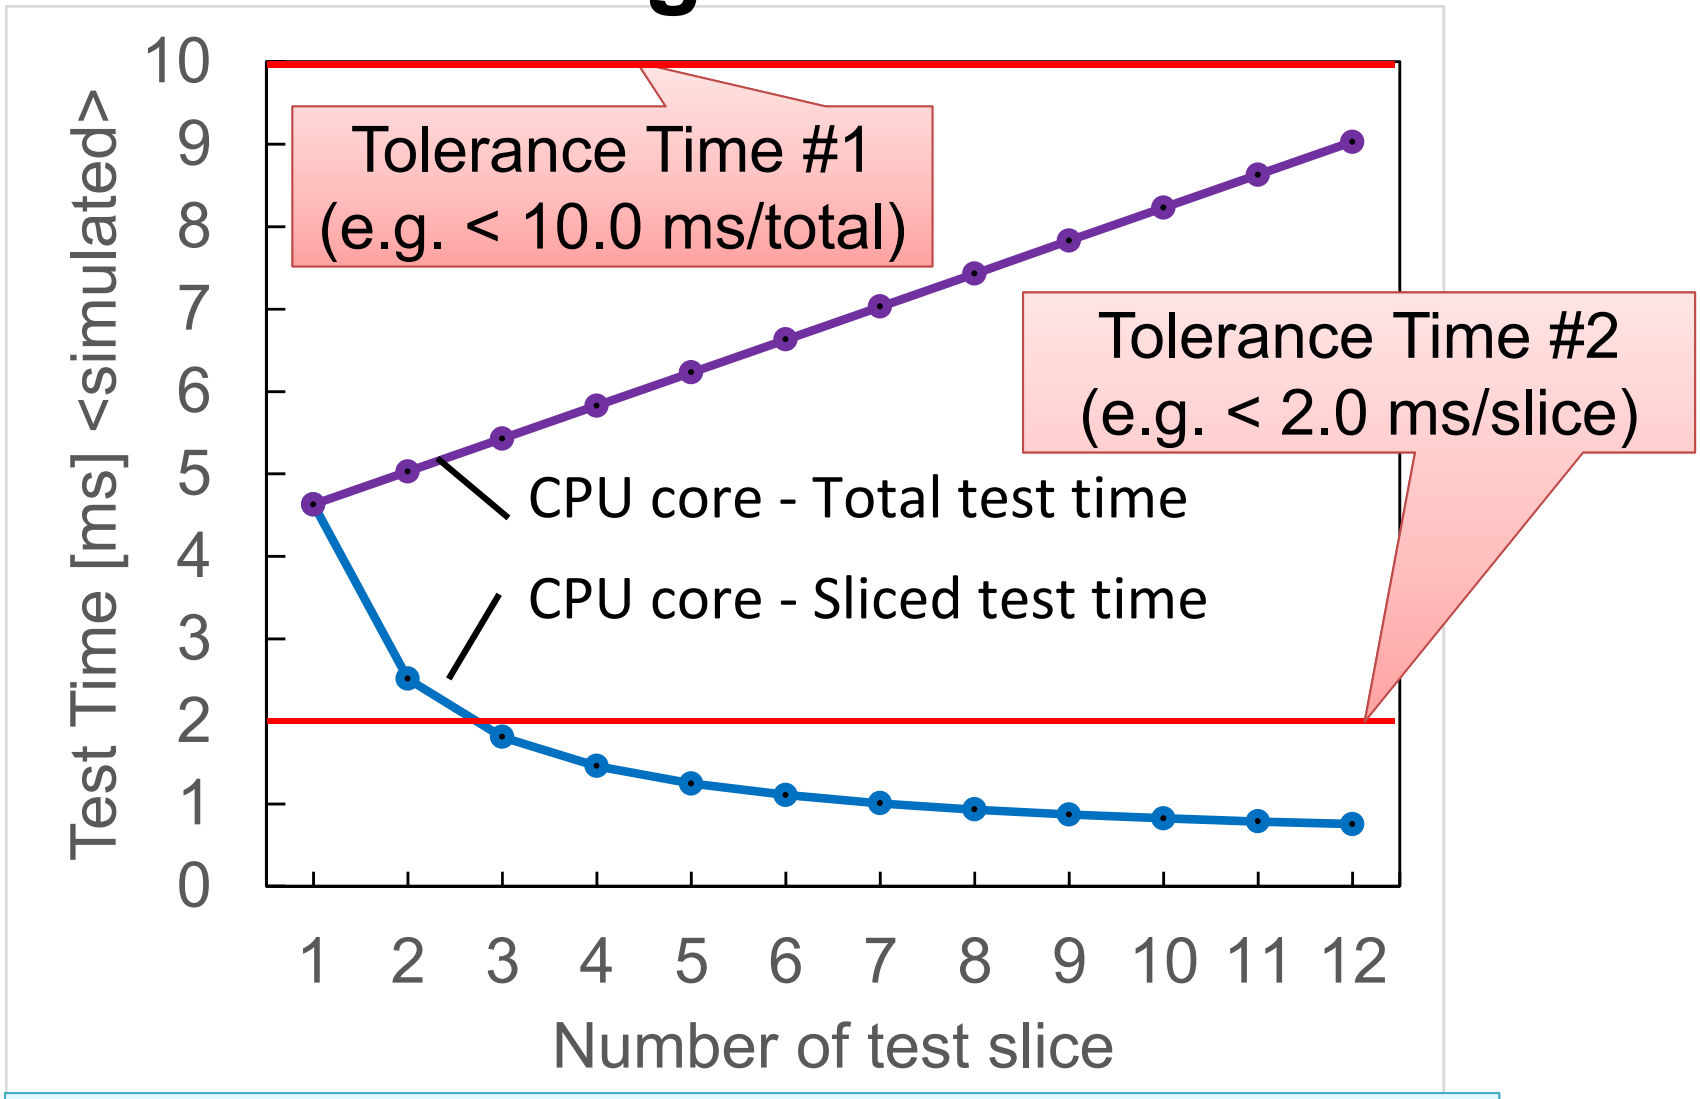

Slicing self-test can satisfy required tolerance time.

Comparison of Safety Mechanisms for Detecting Random H/W Fault

	No measure	Redun-dant HW	Self-test by SW	Self-test w/ HW	This work
Diag-nostic coverage	N/A	**High**	Low	**Medium**	**Medium**
Relative logic size	**1.00**	>2.30	**1.00**	**<1.10**	**<1.10**
Interrupt disconti-nuous time	**N/A**	**N/A**	>100ms	> 10ms	**<2.0ms**

Killer-Droop Monitor with Droop Prediction

- Predicts critical voltage droop in high speed CPU cluster, and prevents fault in case of droop

- Key features are:
  - High speed voltage monitoring
  - Voltage prediction unit
  - 0-cycle delay clock control

Voltage Droop

- Voltage droop is typically caused by change of logic activity from low to high.

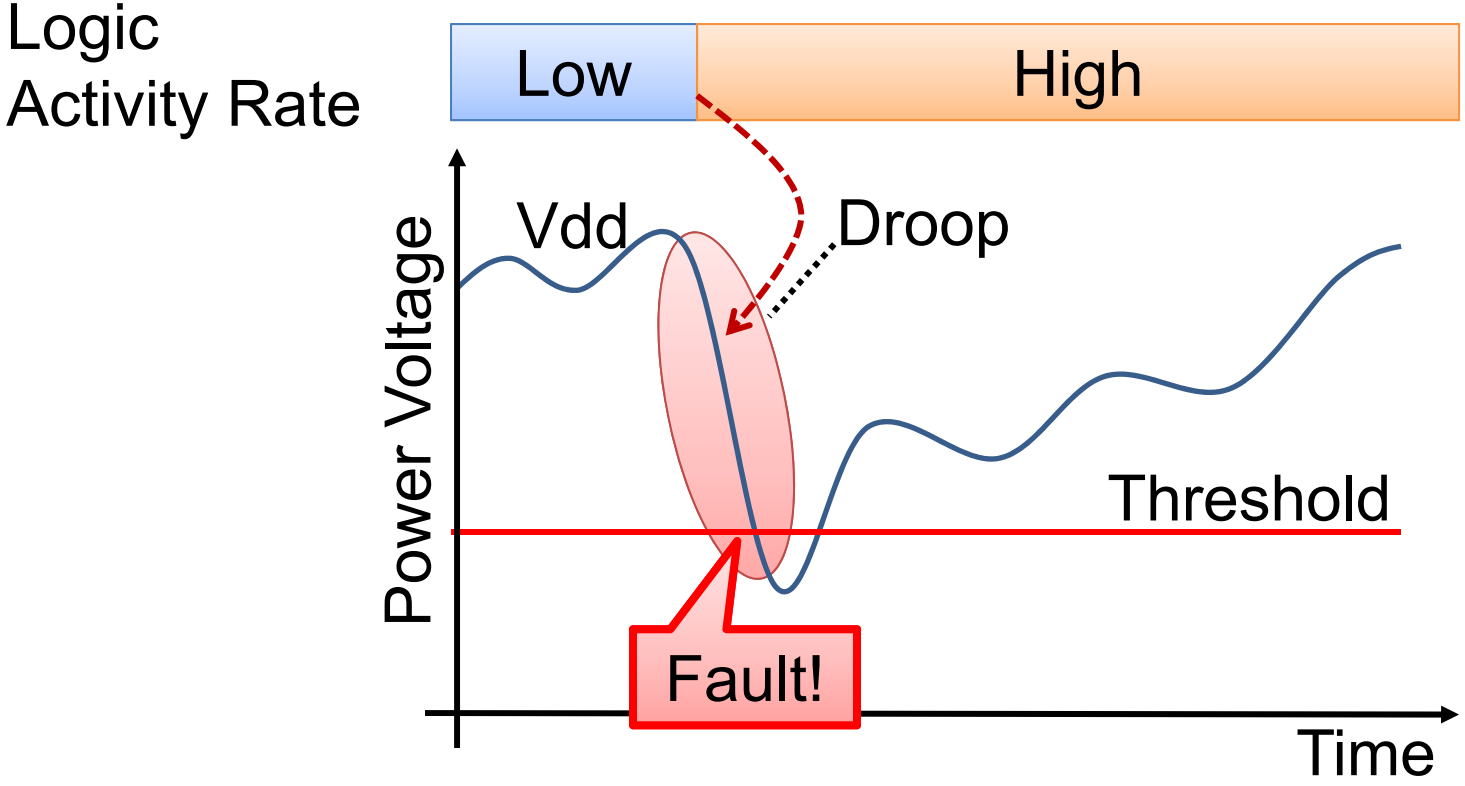

- Design margin cares for the droop, but considering the margin is difficult in low Vdd and higher freq.

Our Strategy against Voltage Droop

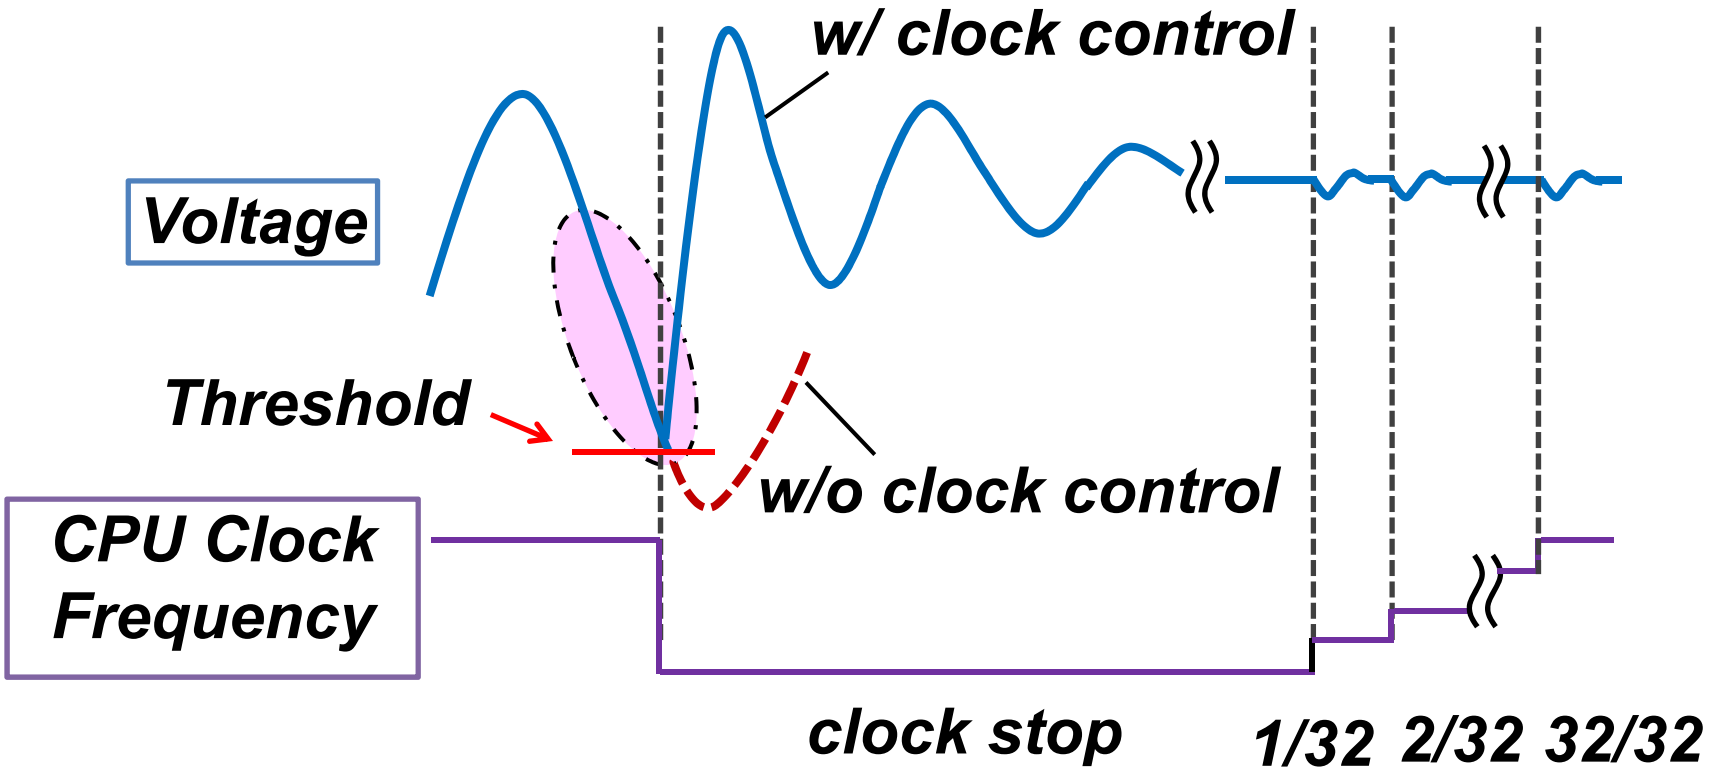

- Introduce a voltage predictor to prevent the threshold violation.
- Clock is stopped until the droop becomes acceptingly small.
- Clock is restarted with 1/32 freq. to avoid large droop.

Block diagram and implementation

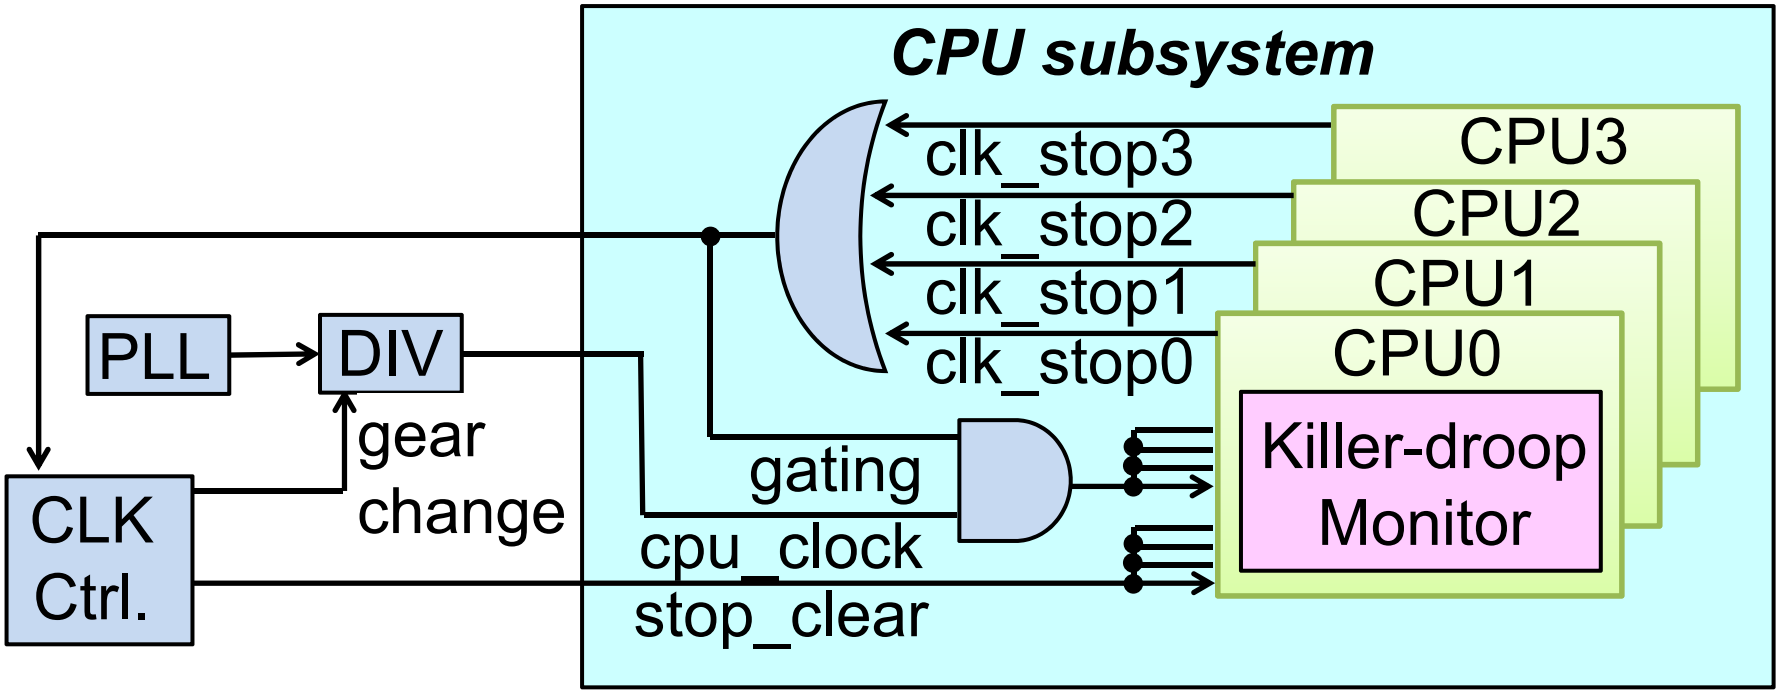

- ‘Killer-droop’ monitor is introduced in each CPU.
- Clock stop is requested if critical voltage droop is predicted.
- Clock controller changes gear of the clock divider, and requests to disable the clock gating.

Droop monitor

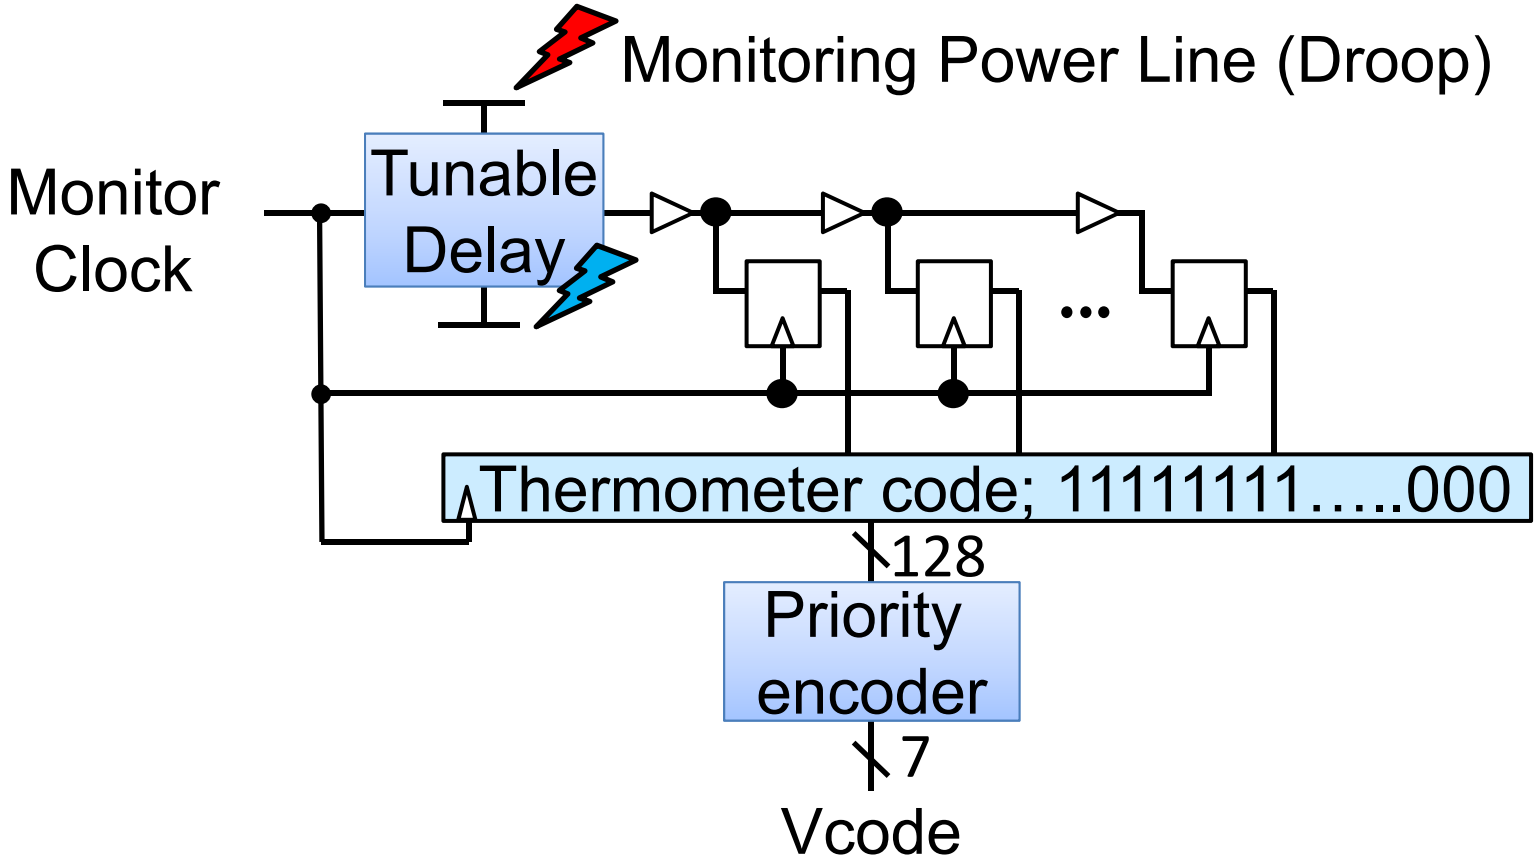

- Based on TDC (Time-to-Digital Converter) architecture.
- Monitored voltage is converted to thermometer code.
- Sampling voltage at up to 2GHz (same as CPU clock).

Droop Prediction

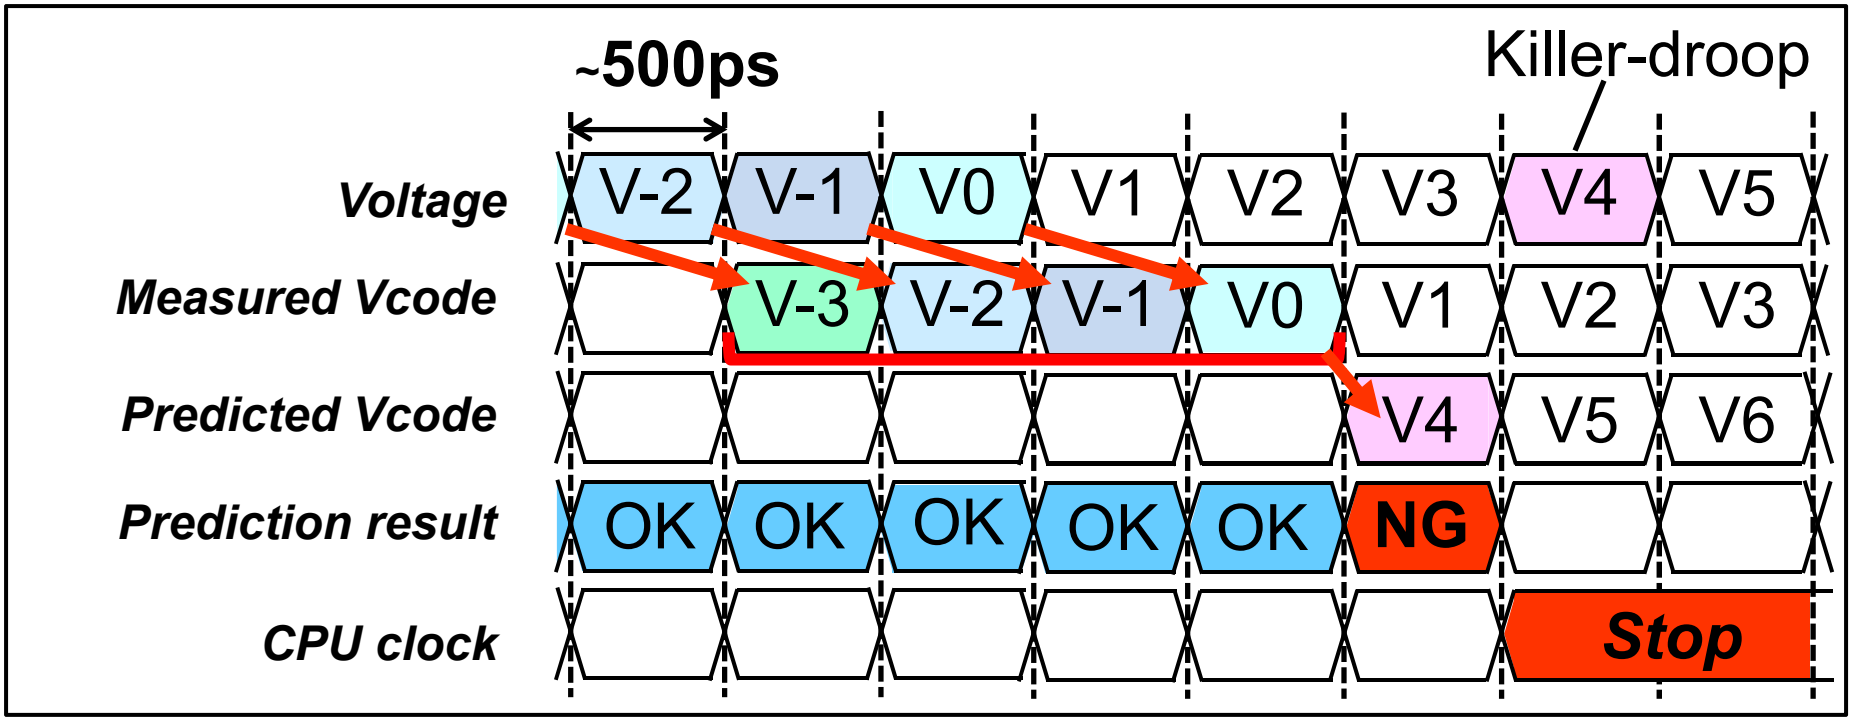

- The monitor takes 3 cycles for prediction. Therefore, it predicts Vcode 4 cycles in advance.
- Droop prediction realizes clock stop with 0-latency when killer-droop is predicted.

Droop Prediction Scheme

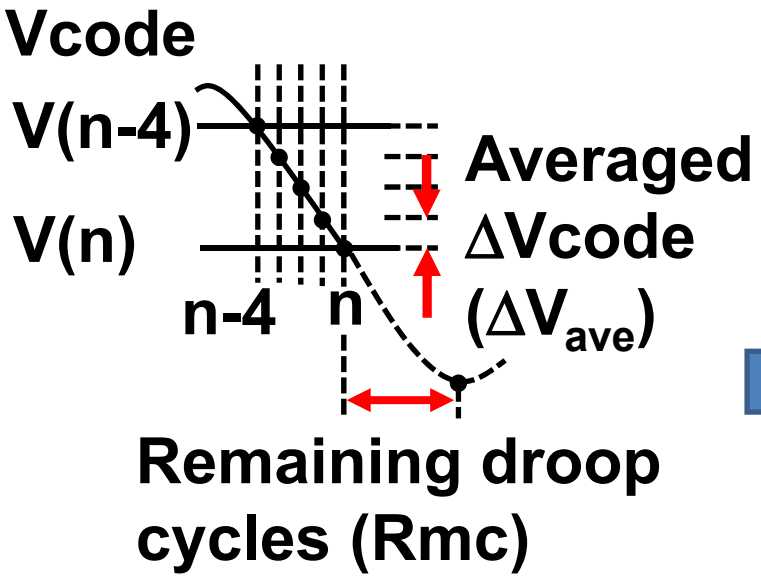

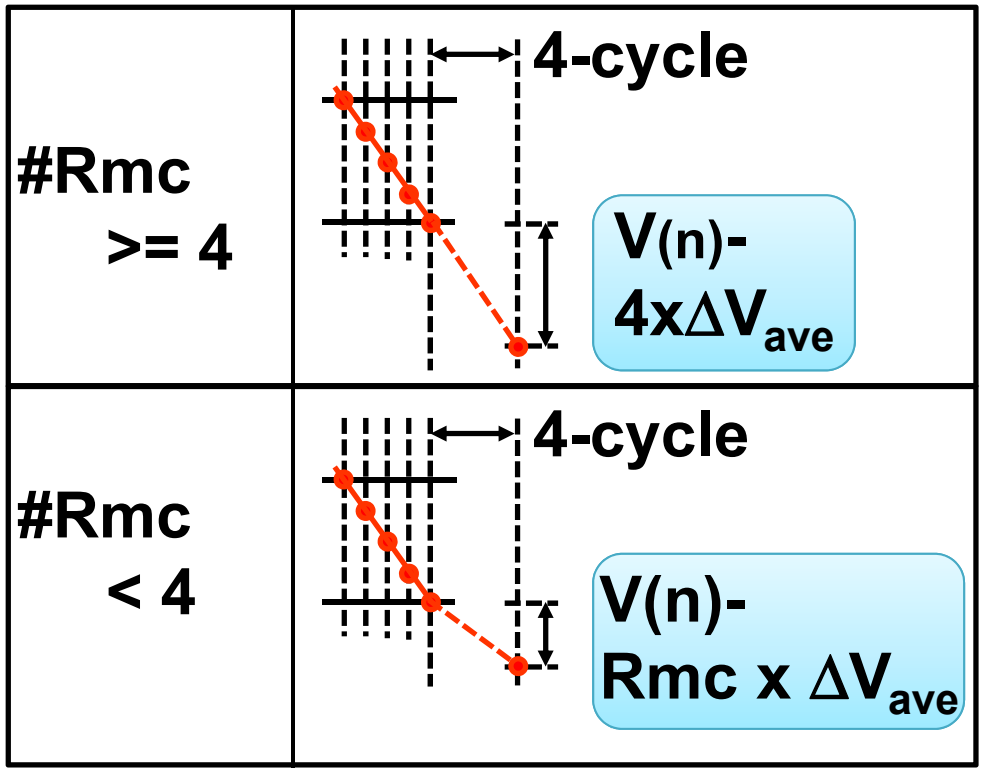

- Uses the least 4cycle ΔVcode history and remaining droop cycles.
- Remaining droop cycles is calculated by Vcode history and anti-resonant frequency.

Droop Prediction Result

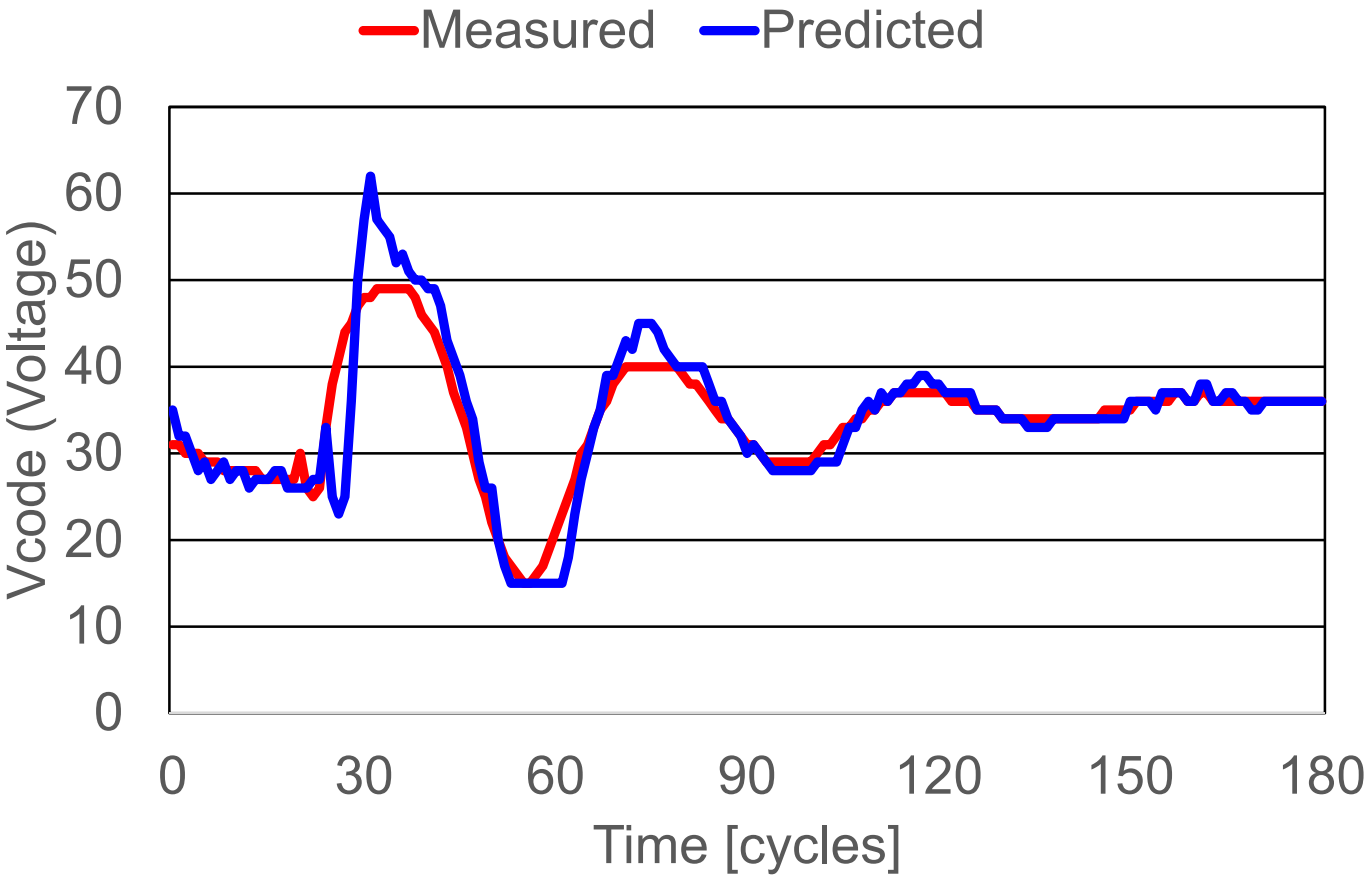

- Minimum droop voltage is predicted with good accuracy.
- Predicted minimum Vcode is equal or lower than measured.
- Droop max voltage has some differences from the prediction, but maximum voltage is not used for clock stop.

Shmoo plot

w/o Killer-droop monitor

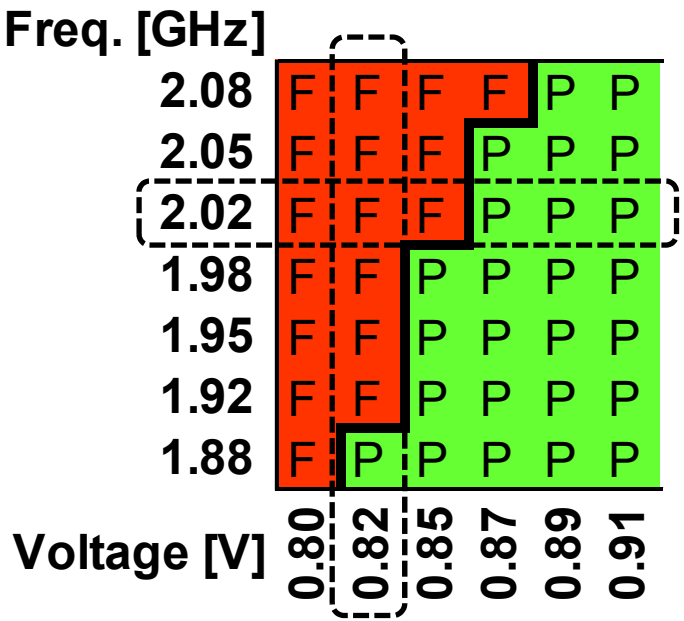

2.02GHz@0.87V

1.88GHz@0.82V

w/ Killer-droop monitor

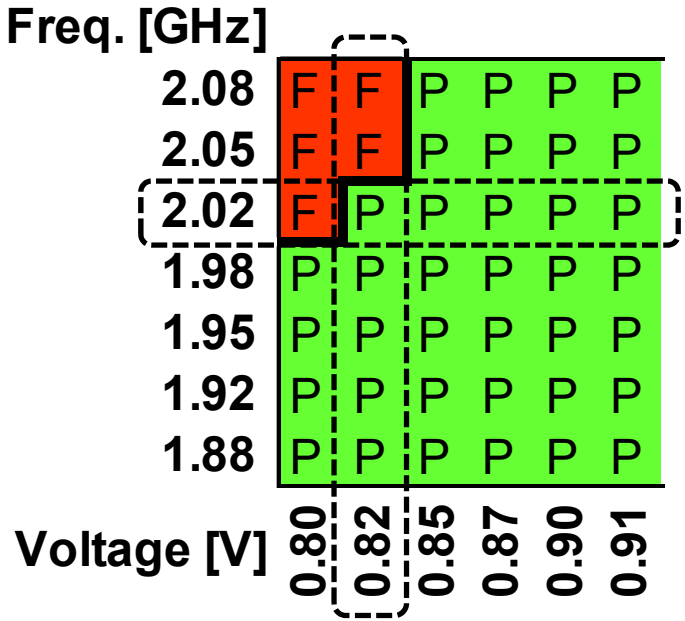

2.02GHz@0.82V

Minimum Vdd improved by: 50mV@2.02GHz

Frequency improved by: 140MHz@0.82V

Comparison with other works

	[2]	[3]	[4]	This work
Technology	28nm	28nm	16nm	16nm
Application	Mobile	PC	Mobile	Automotive
Droop Monitor	Ring oscillator	DLL	Replica path delay	TDC
Reaction Mechanism	Frequency down	Clock pulse stretch	Inserting clock delay	Clock stop
Monitor Interval	50ns	**1-cycle**	**1-cycle**	**1-cycle**
Droop Prediction	n/a	n/a	n/a	**Yes**
Reaction Latency	100ns	2-cycle	1.5-cycle	**0-cycle**

Conclusion

We designed a 16nm FinFET heterogeneous nona-core SoC supporting ISO26262 ASIL B.

- High flexibility run-time self test minimizes discontinuous time with enabling to satisfy required diagnostic coverage for ASIL B.

- Killer-droop monitor predicts critical voltage droop and improves minimum Vdd by 50mV at 2.02GHz.

Demo session

We will show a demo at paper 4.4 booth (Not 4.5)

- Run-time self test on GPU
- Killer-droop monitor

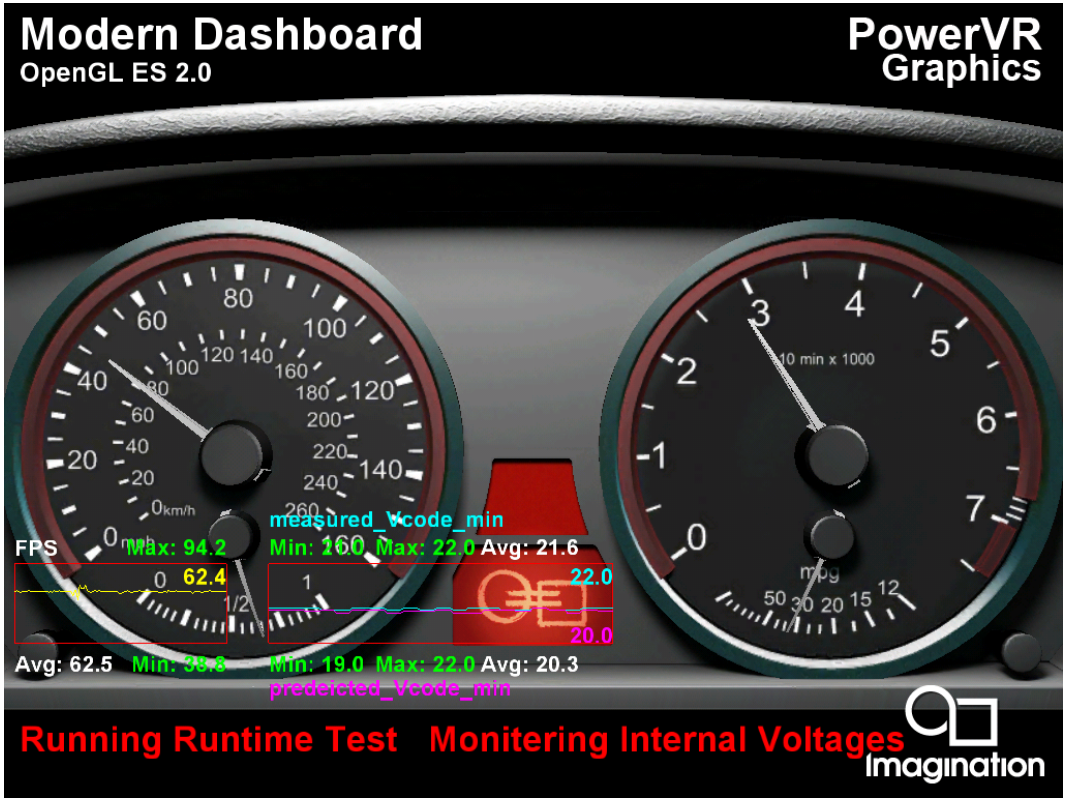

A 65nm CMOS 6.4-to-29.2pJ/FLOP@0.8V Shared Logarithmic Floating Point Unit for Acceleration of Nonlinear Function Kernels in a Tightly Coupled Processor Cluster

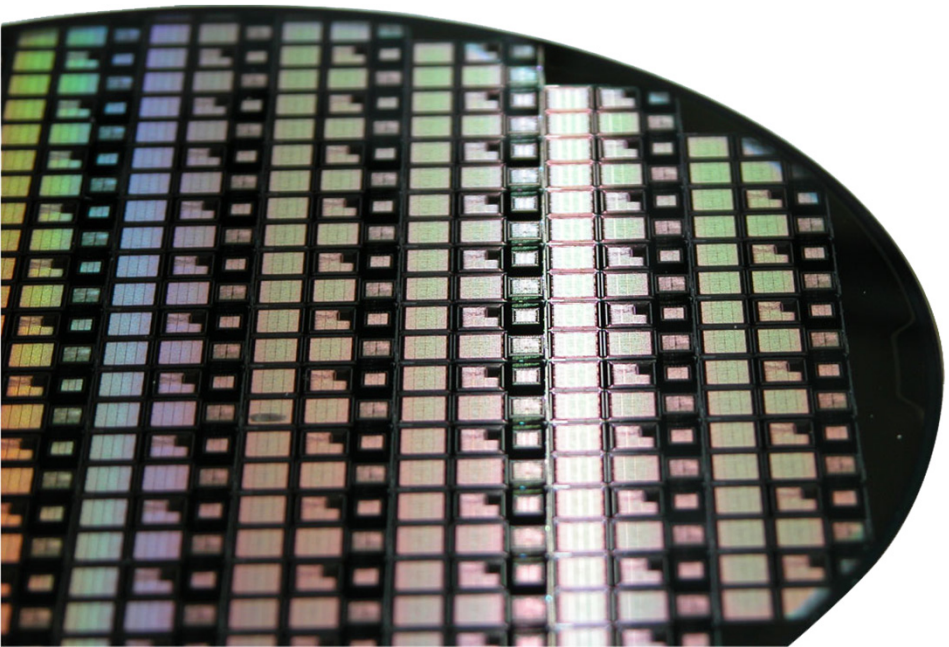

ISSCC 2016
San Francisco
1st - 4th Feb., 2016

Michael Gautschi[1]
Michael Schaffner[1]
Frank K. Gürkaynak[1]
Prof. Luca Benini[1,2]

[1]Integrated Systems Laboratory

[2] Università di Bologna

Advanced processing in IoT

Sense **Analyze and Classify** **Transmit**

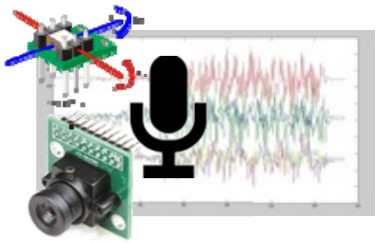

ULP Processing System

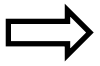

Complex preprocessing close to sensor, e.g.:

Feature extraction, regression, classification,

compression, sensor fusion

Arithmetic with HDR desirable

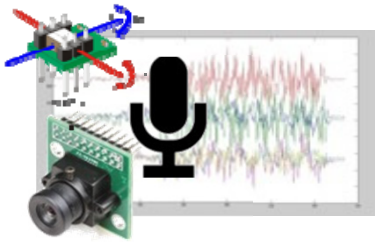

ULP Processing System

100 µW - 2 mW

Fixed-Point
1 - 10 mW

Idle: ~1µW
Active: ~ 50mW

- Fixed-point: labor intensive, error-prone, quality losses

Arithmetic with HDR desirable

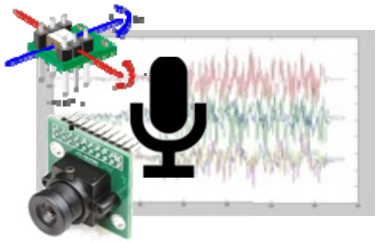

ULP Processing System

100 µW - 2 mW

HDR Arithmetic
1 - 10 mW

Idle: ~1µW
Active: ~ 50mW

- Fixed-point: labor intensive, error-prone, quality losses
- Energy-efficient, low-cost HDR arithmetic desirable

Logarithmic Number System (LNS)

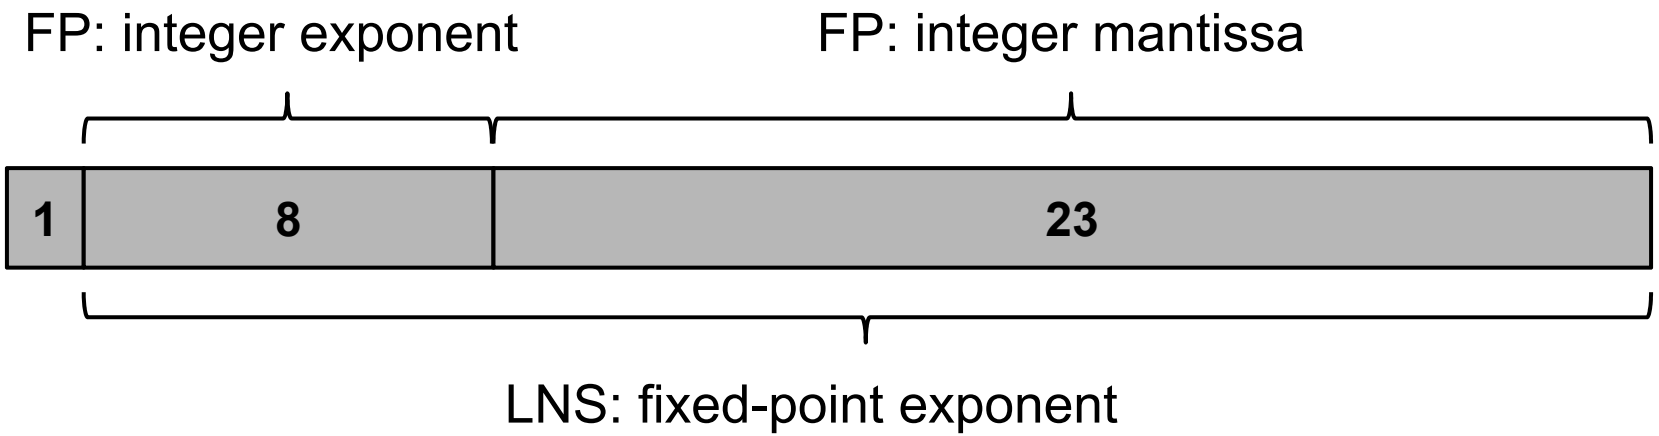

- Alternative for standard Floating Point (FP)
- A = sign * $2^{(\text{fixed point number})}$
- Nonlinear ADD, SUB, I2F, F2I
  - function interpolator → large LNS FPU (LNU)

Efficient LNS operations

Multiplication/ Division:

- $C = A \; */ \; B$ with $A = 2^a, \; B = 2^b, \; C = 2^c$
- $c = \log_2(2^a \; */ \; 2^b) = \log_2(2^{a \pm b}) = a \pm b$

Square-roots:

- $C = sqrt(A)$ with $A = 2^a, \; C = 2^c$
- $c = \log_2(sqrt(2^a)) = \log_2(2^{0.5a}) = 0.5a = a >> 1$

Simple integer operations!

Private Floating Point Units

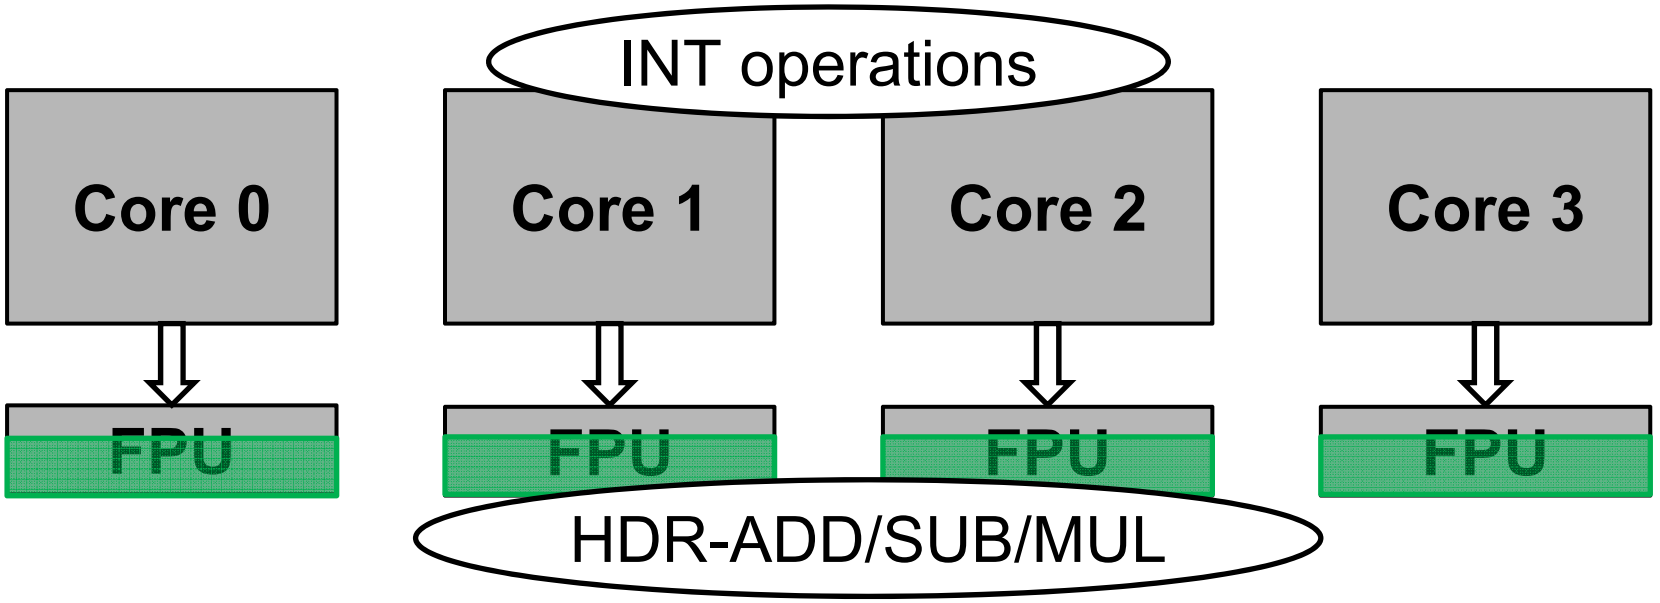

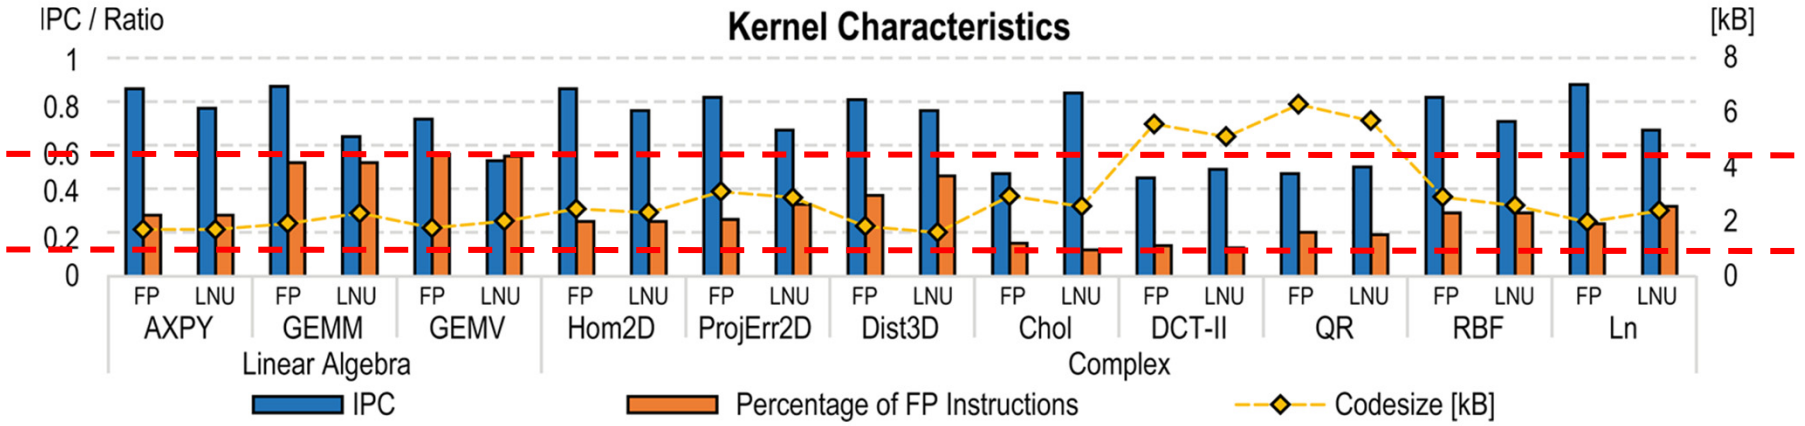

Logartihmic Number System (LNU)

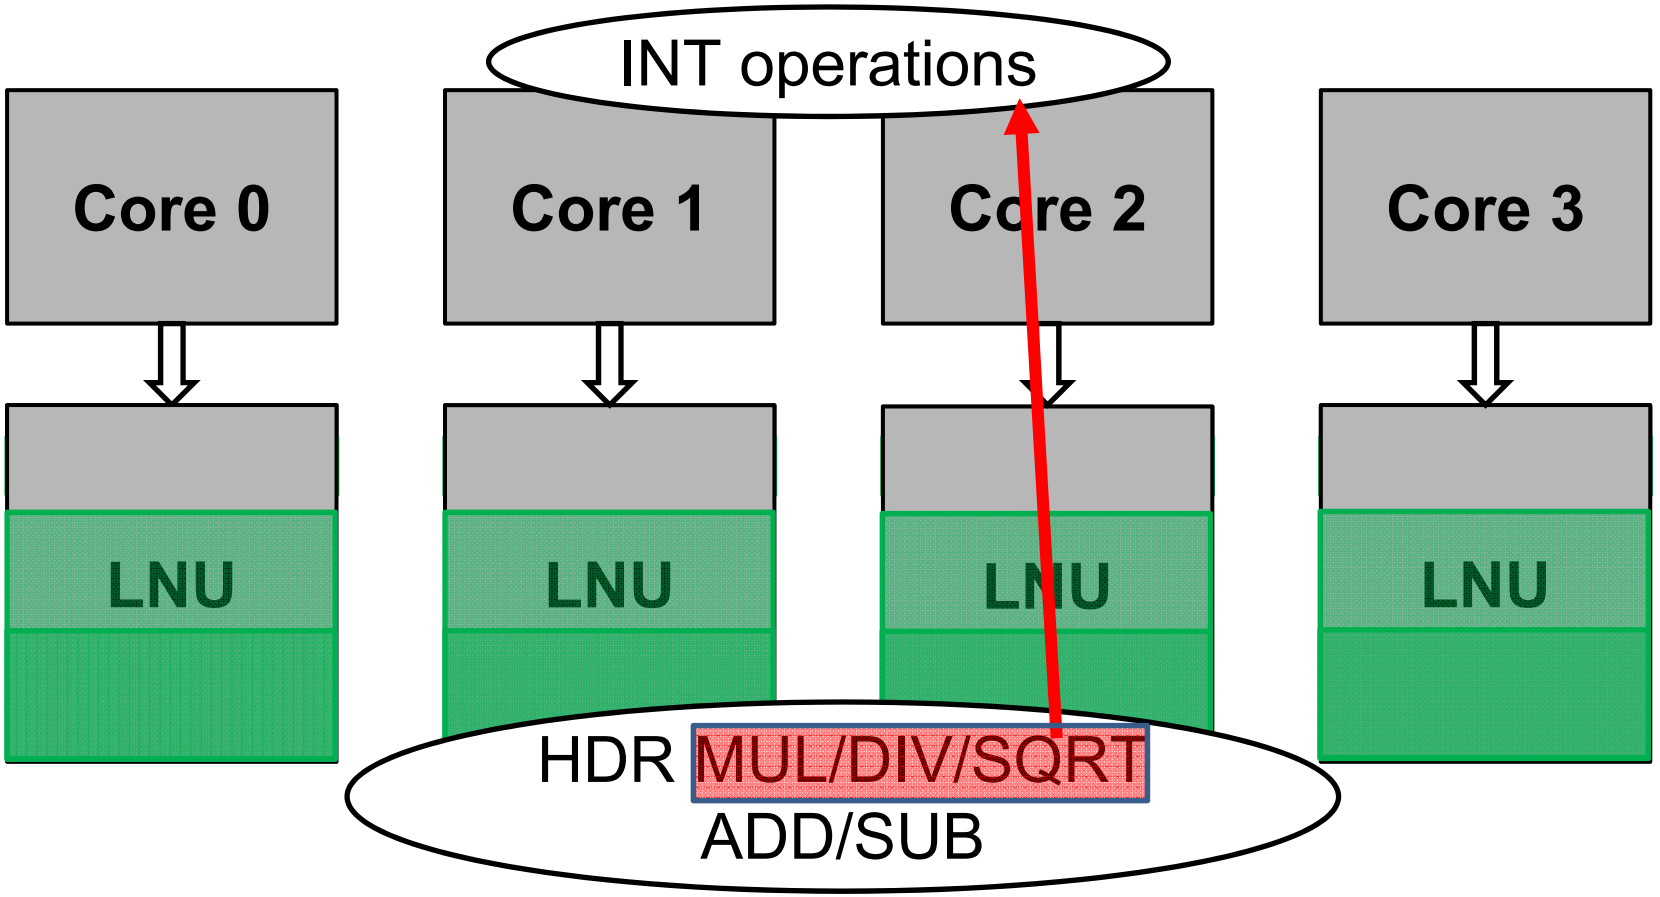

- Area: 1 LNU ≈ 5 × standard IEEE compliant FPU (no DIV)
- Poor LNU utilization ~ 0.2

Shared LNU

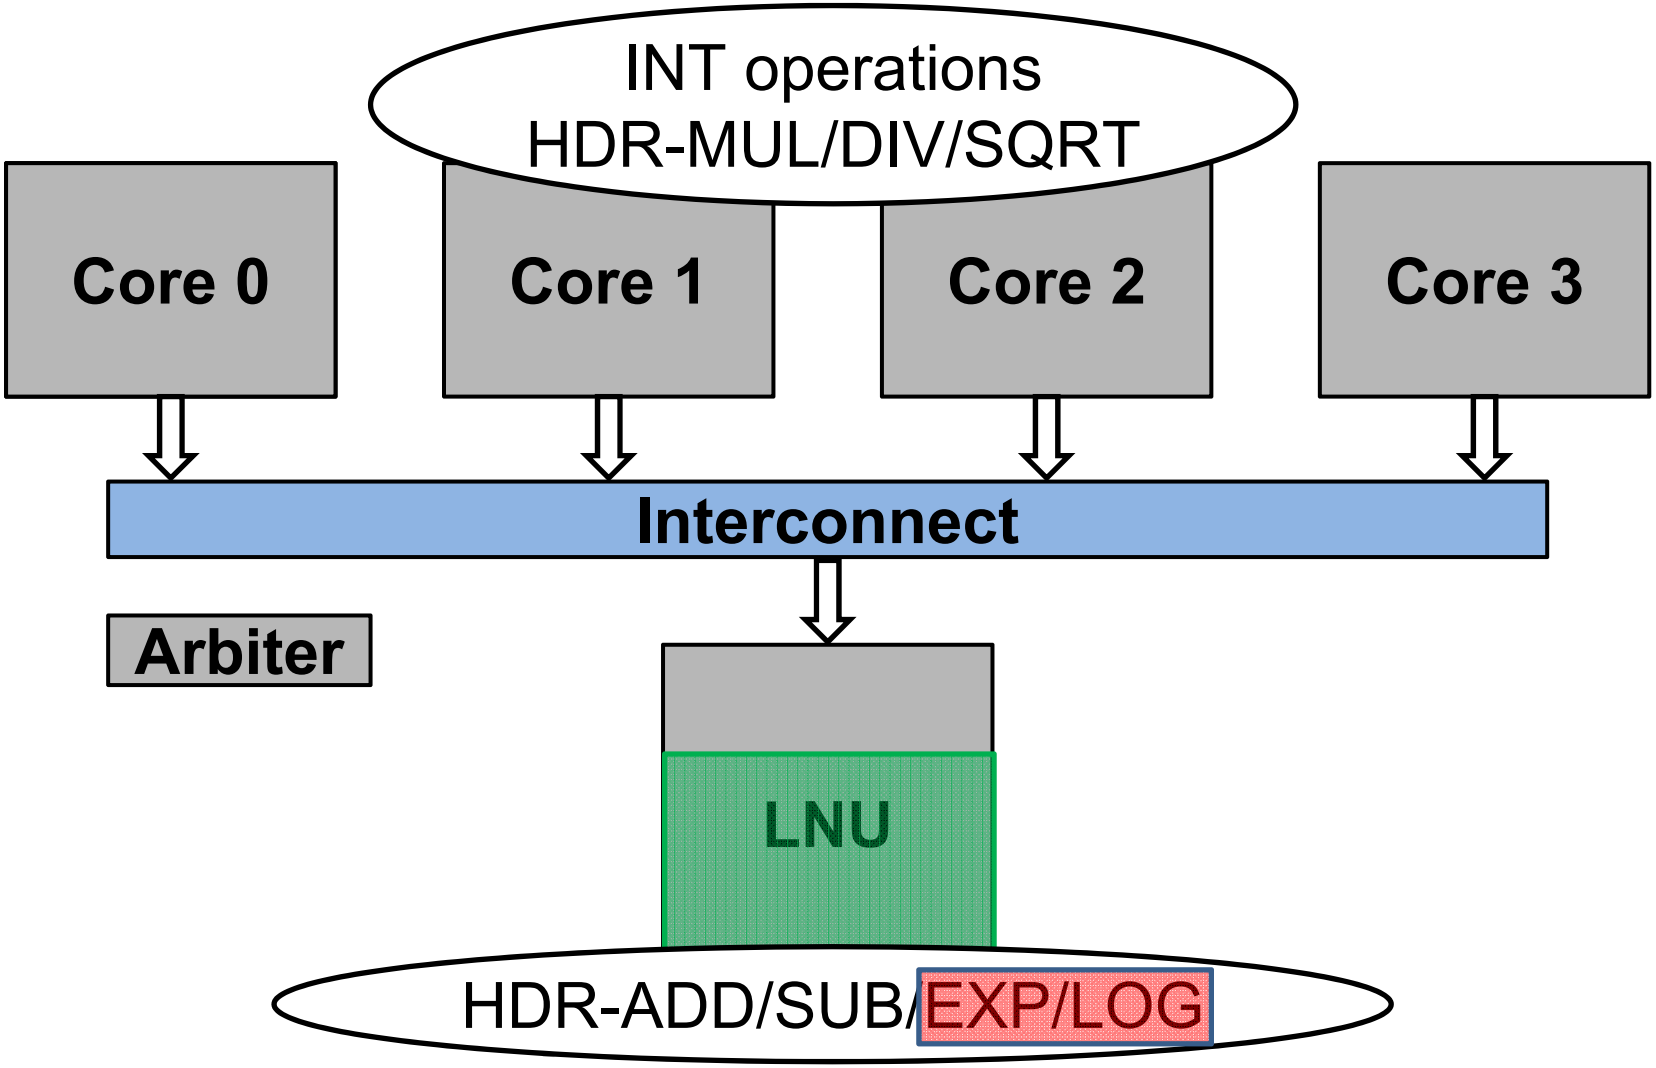

Contributions

- Design of **35% smaller** LNU
- **More functionality** than previous designs [1,2]
  (EXP,LOG, I2F, F2I)
- **First multi-core** chip with a **shared** LNU
- **Energy efficiency measurements** in 65nm CMOS

[1] J.N. Coleman et al. "The European Logarithmic Microprocessor" IEEE TC, 2008
[2] R.C. Ismail et al. "ROM-less LNS" IEEE ARITH, 2011

Outline

- LNS Additions and Subtraction
- Architecture of the Logarithmic Number Unit
- Multicore Hardware Platform
- Results
- Conclusion
- Q&A

Problematic LNS Additions/Subtractions

- $C = A \pm B$ with $A = 2^a,\ B = 2^b,\ C = 2^c$
- Easy case (ADD): $c = \log_2(2^a + 2^b) = \max(a,b) + f^+(-|a-b|)$
- Hard case (SUB): $c = \log_2(2^a - 2^b) = \max(a,b) + f^-(-|a-b|)$

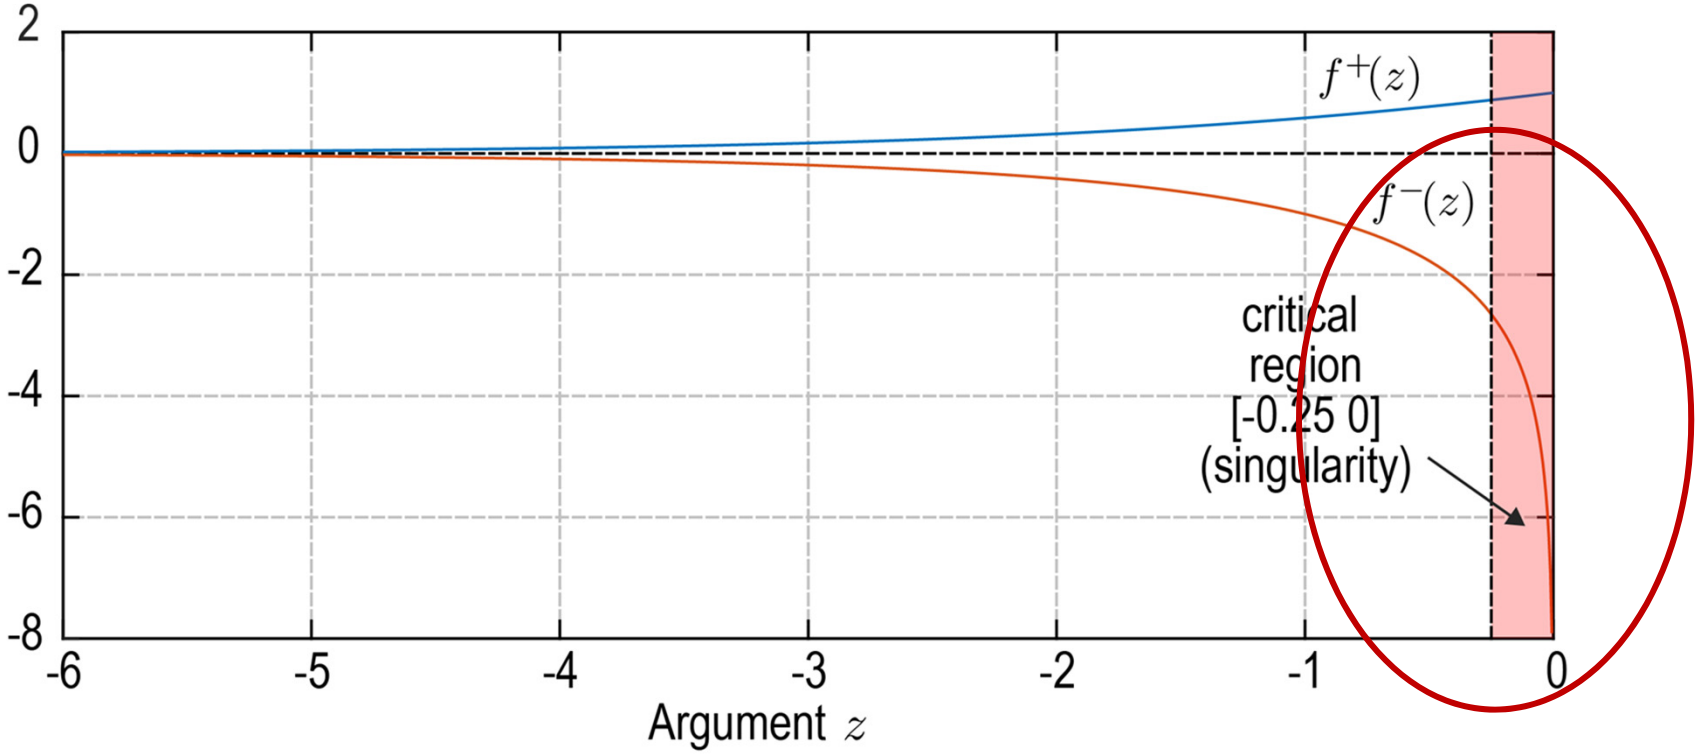

Critical Region Decomposition

$$c = \max(a,b) + f^{-}(z)$$

$$c = \max(a,b) + \log_2((1-2^z) / (1+2^z)) + f^{+}(z)$$

co-transformation(z)

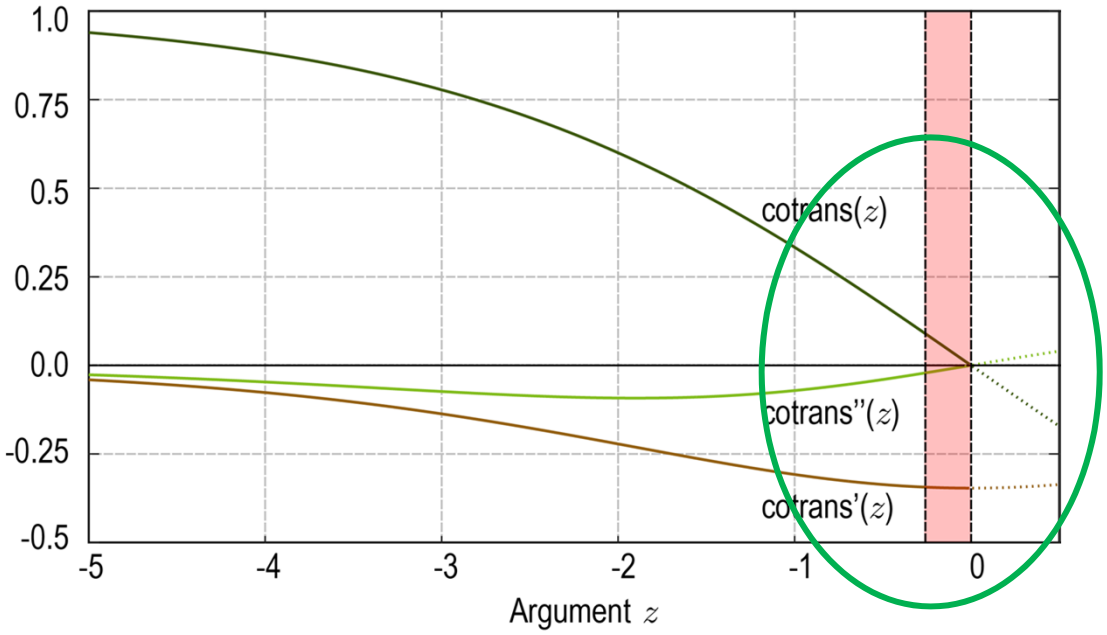

Function Approximation

- 32b single precision → **max error 0.5 ulp**
- 1 ulp = 2^-23
- Different methods to approximate f^+, f^- , *e.g.*,

1) LUT only (very large!)
2) High order polynomial
3) LUT + piecewise poly (tradeoff)

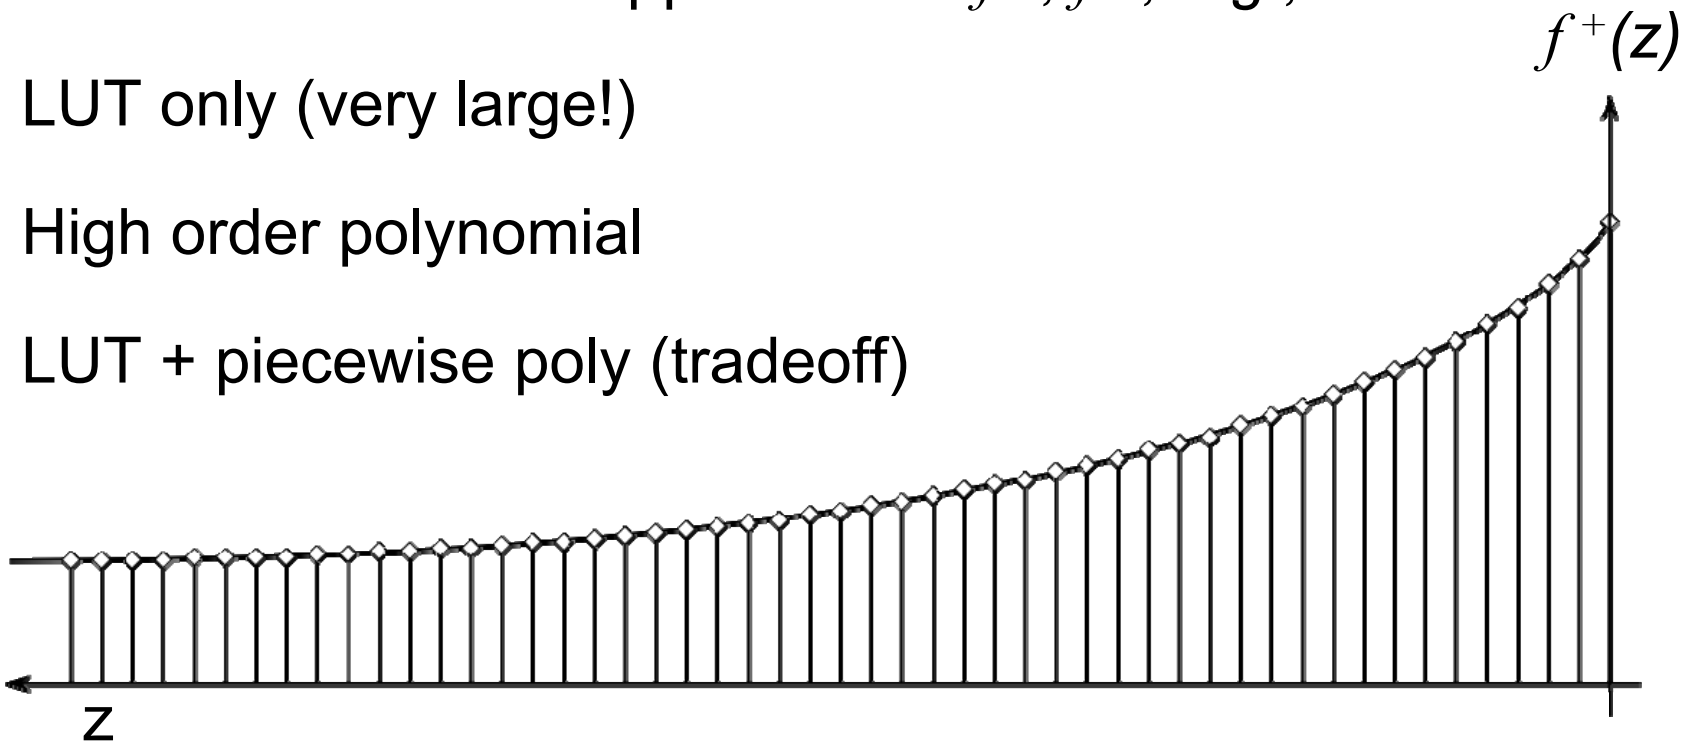

Interpolation Order

	LUT based	1st order	2nd order	3rd order
LUT size	--	-	+	++
Interpolator delay/size		++	+	-
AT tradeoff	-	-	+	-

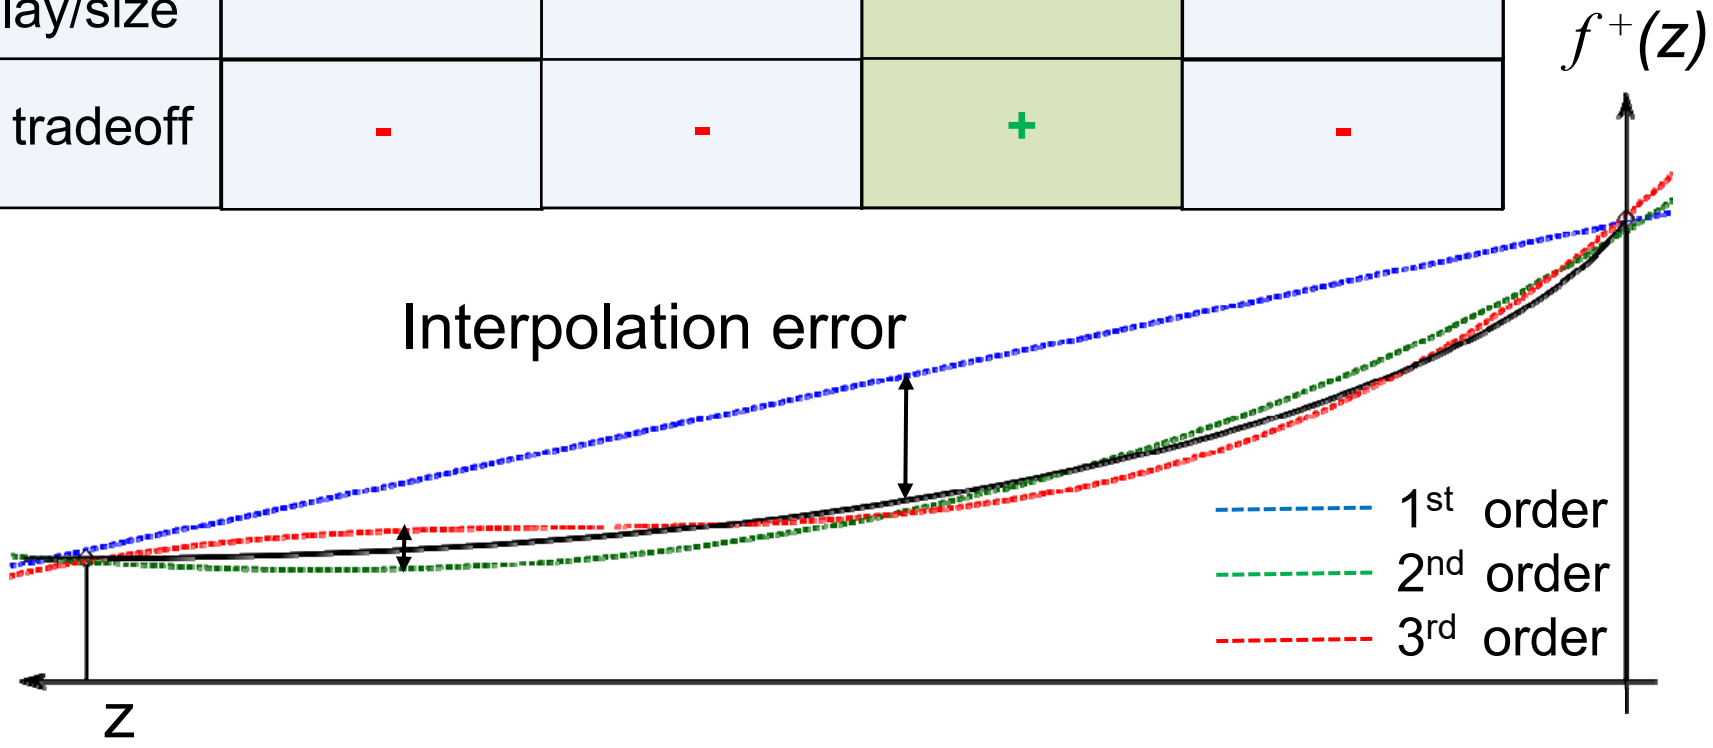

Piecewise Polynomials

	LUT based	1st order	2nd order	3rd order
LUT size	--	-	+	++
Interpolator delay/size		++	+	-
AT tradeoff	-	-	+	-

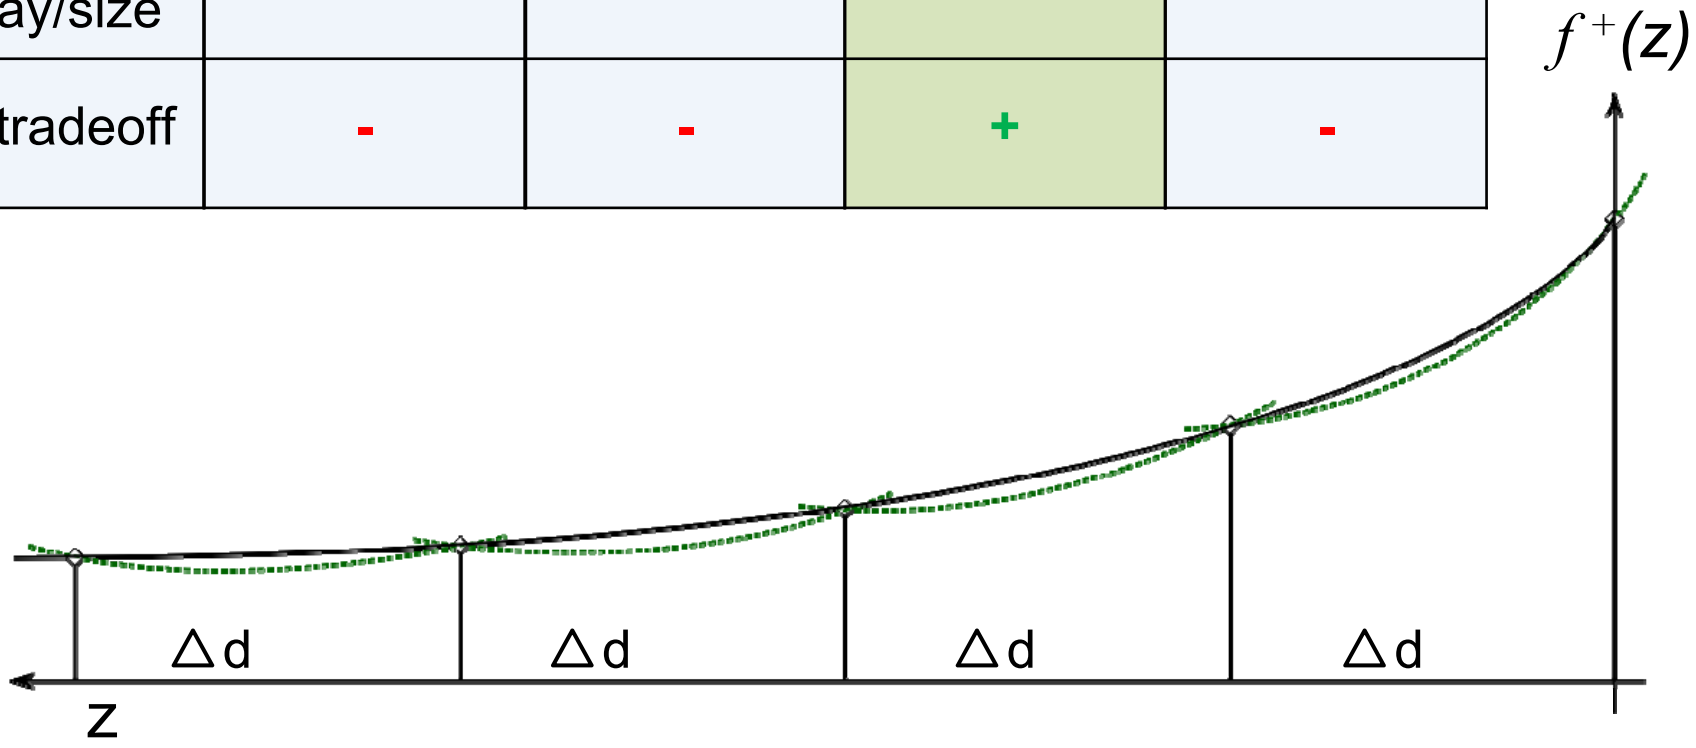

LUT Segmentation

	LUT based	1st order	2nd order	3rd order
LUT size	--	-	++	++
Interpolator delay/size		++	+	-
AT tradeoff	-	-	+	-

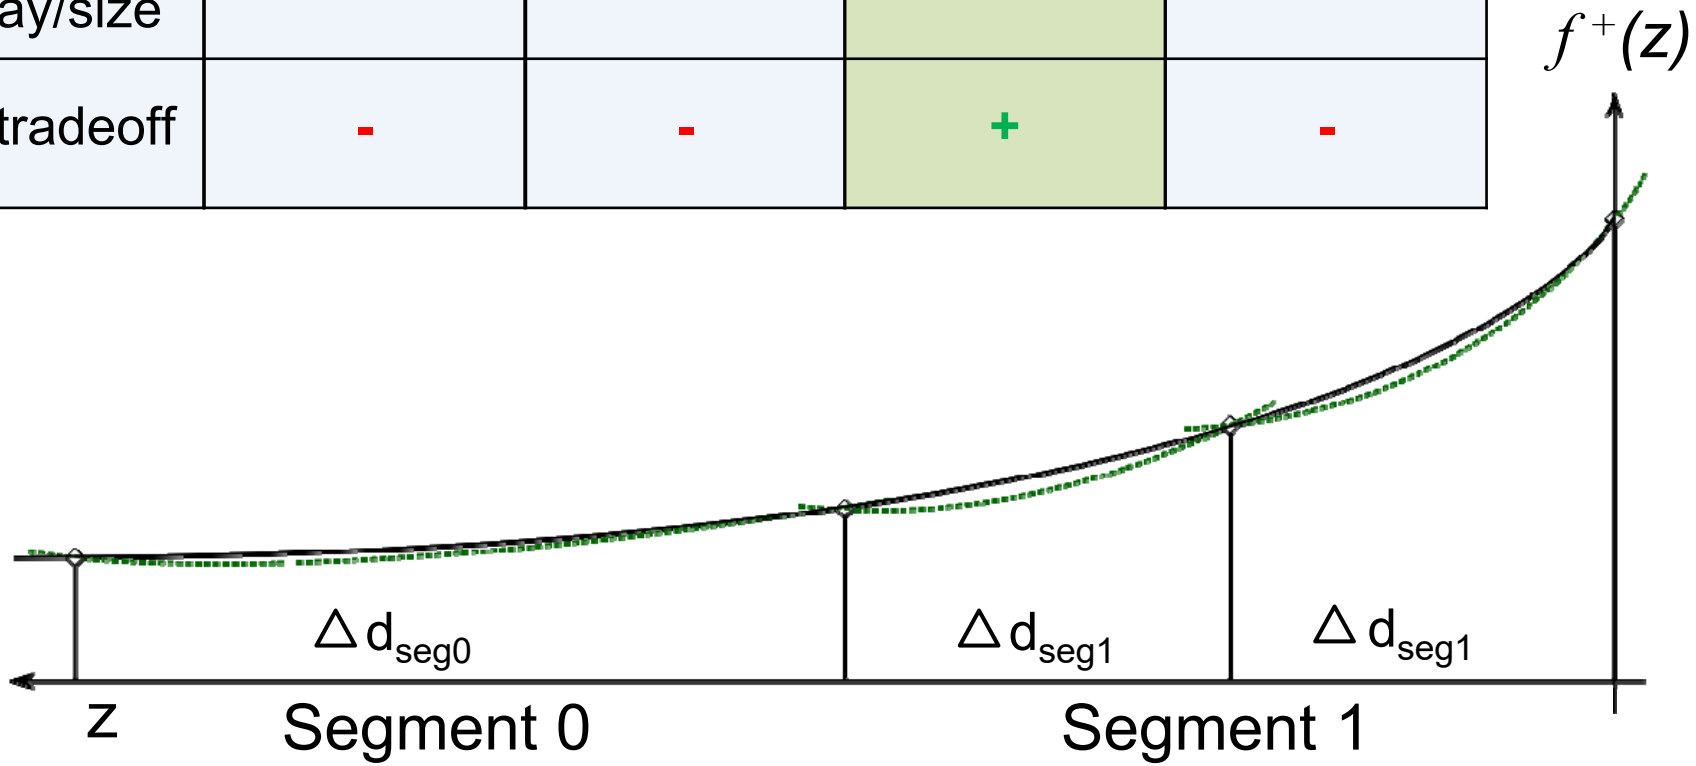

Logarithmic Number Unit (LNU)

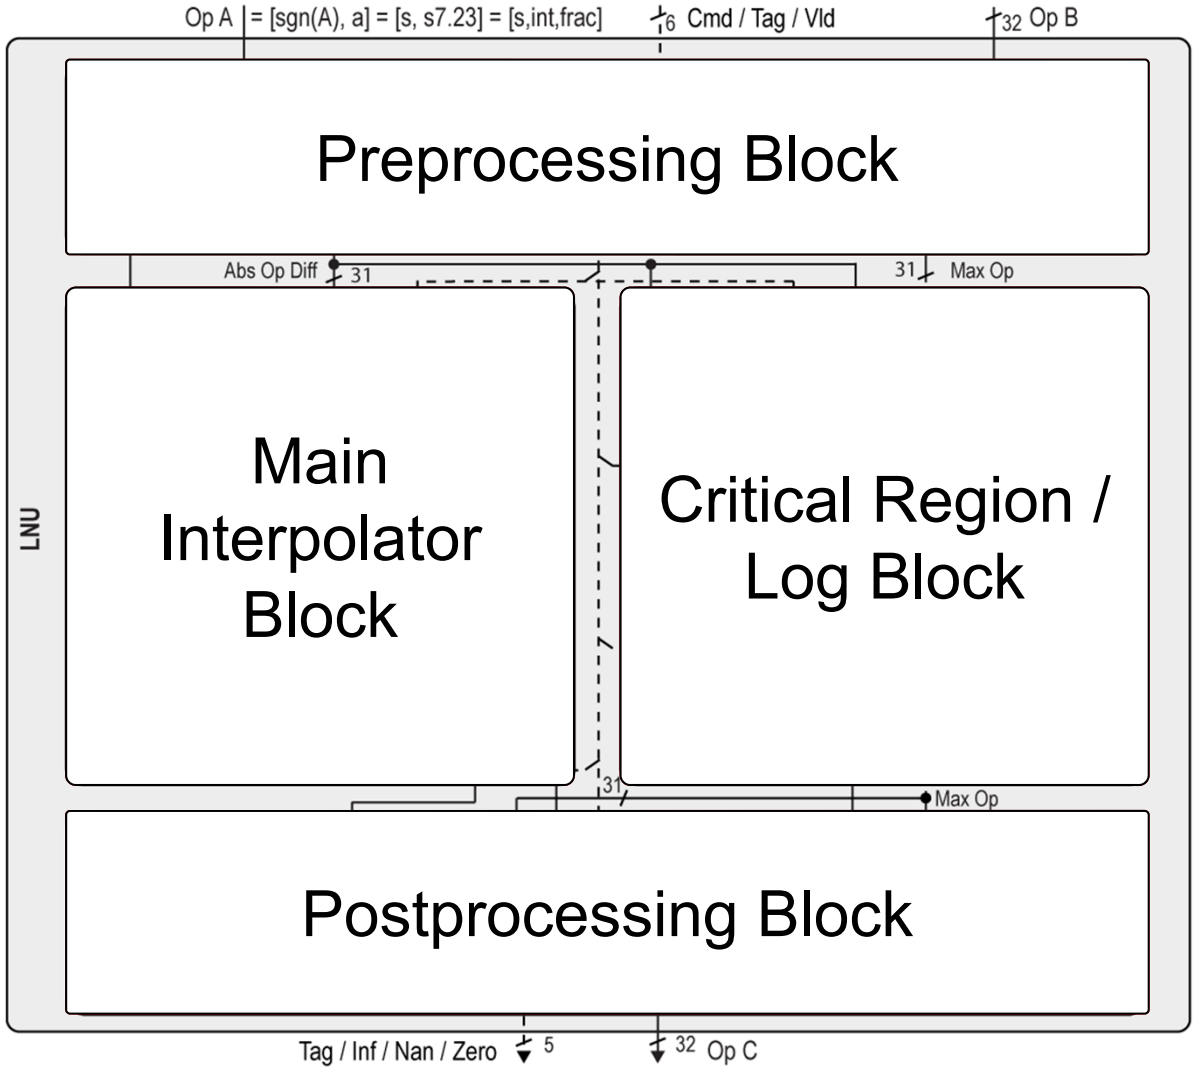

LNS SUB (critical):

$$c = \max(a,b) + \log_2(co\text{-}trans(z)) + f^{+}(z)$$

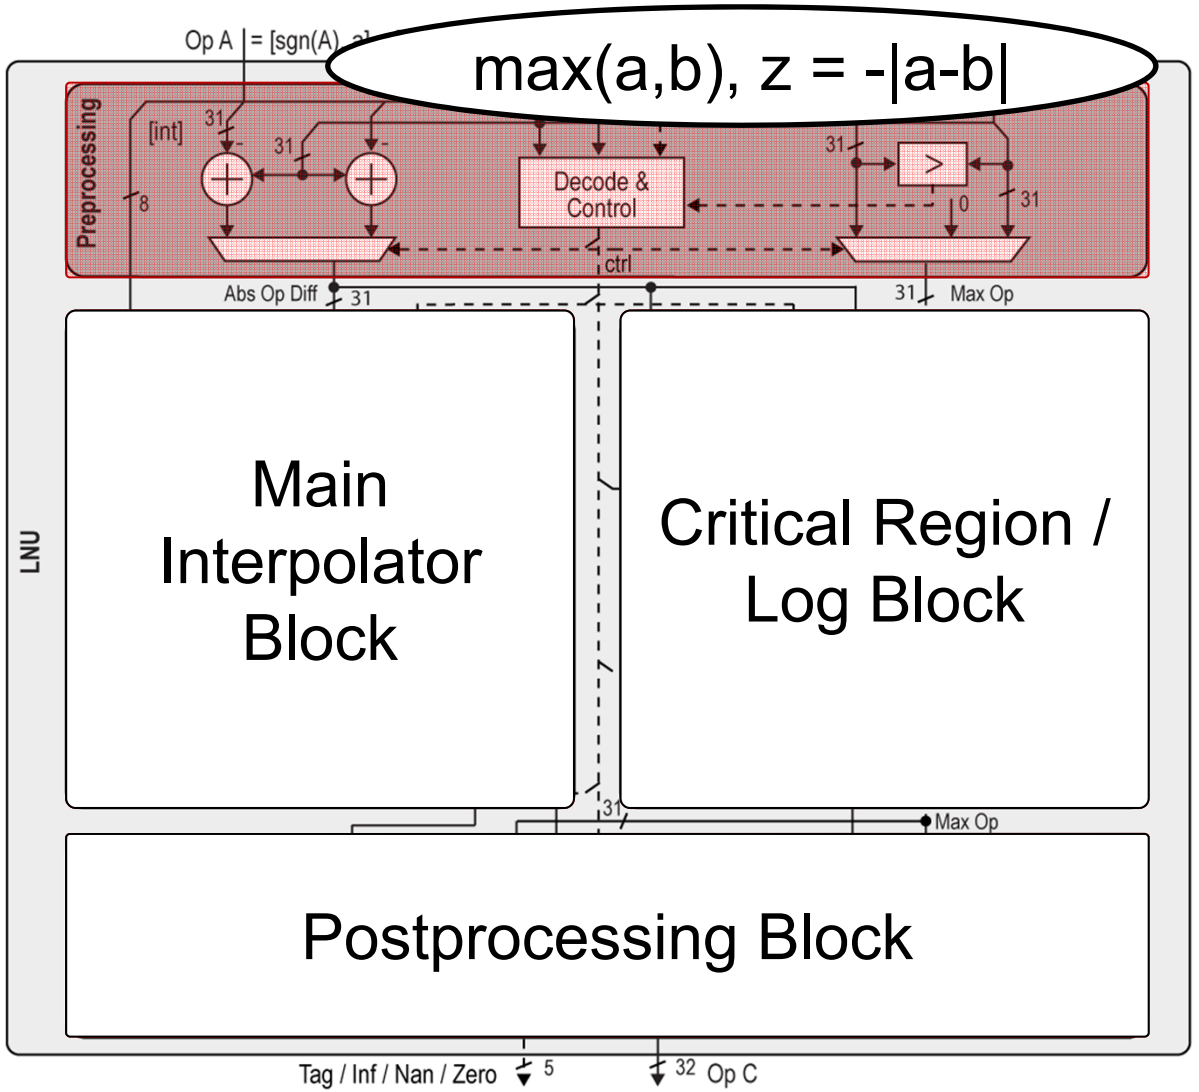

LNS SUB (critical):

$$c = \max(a,b) + \log_2(co\text{-}trans(z)) + f^+(z)$$

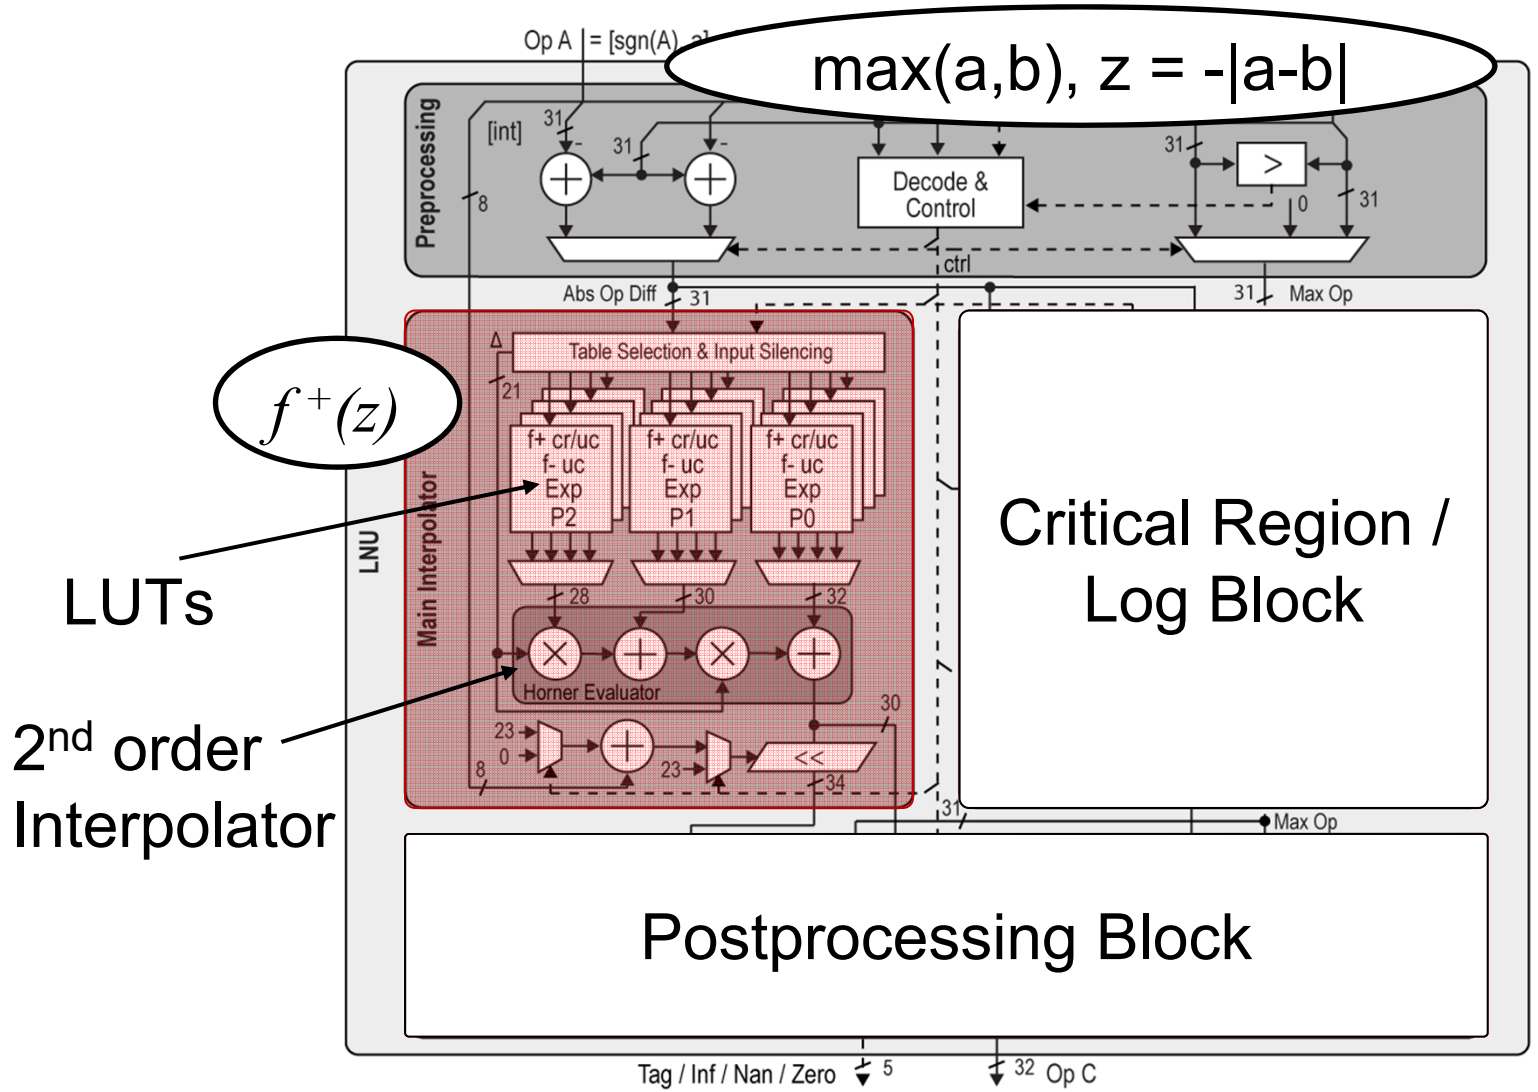

LNS SUB (critical):

$$c = \max(a,b) + \log_2(co\text{-}trans(z)) + f^+(z)$$

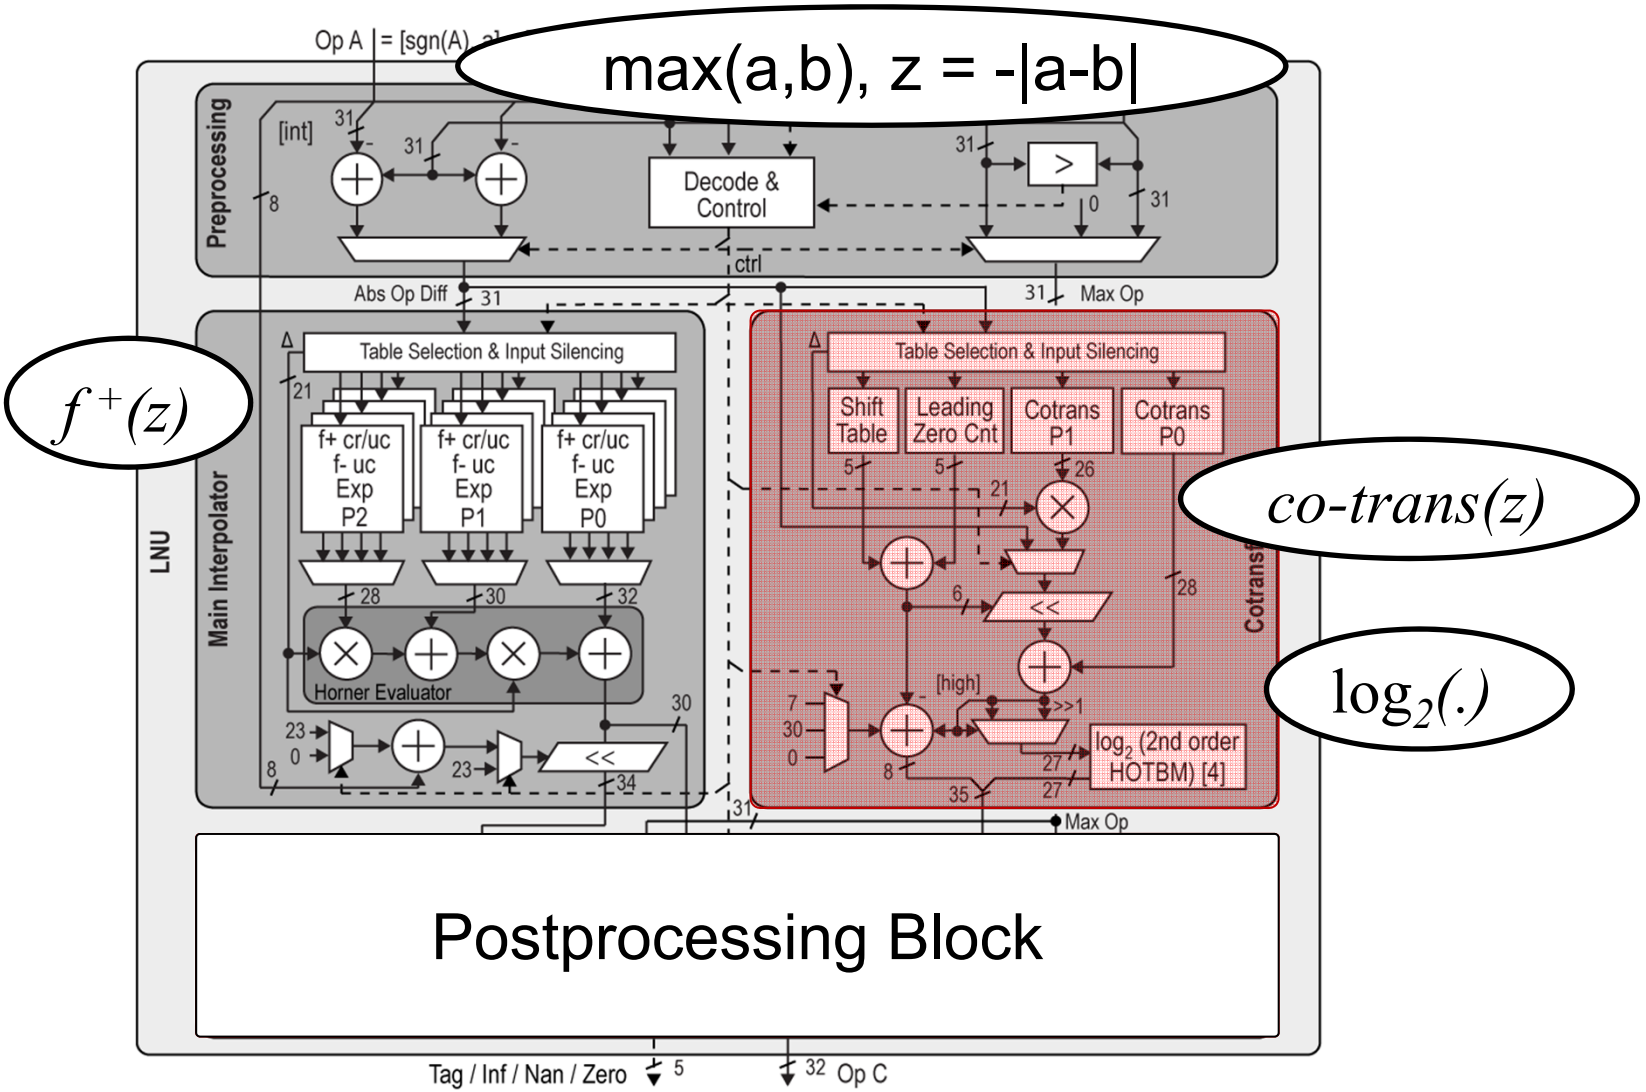

[4] Detrey, J. et. al. "Table-based polynomials for fast hardware function evaluation." ASAP 2005

LNS SUB (critical):

$$c = \max(a,b) + \log_2(co\text{-}trans(z)) + f^+(z)$$

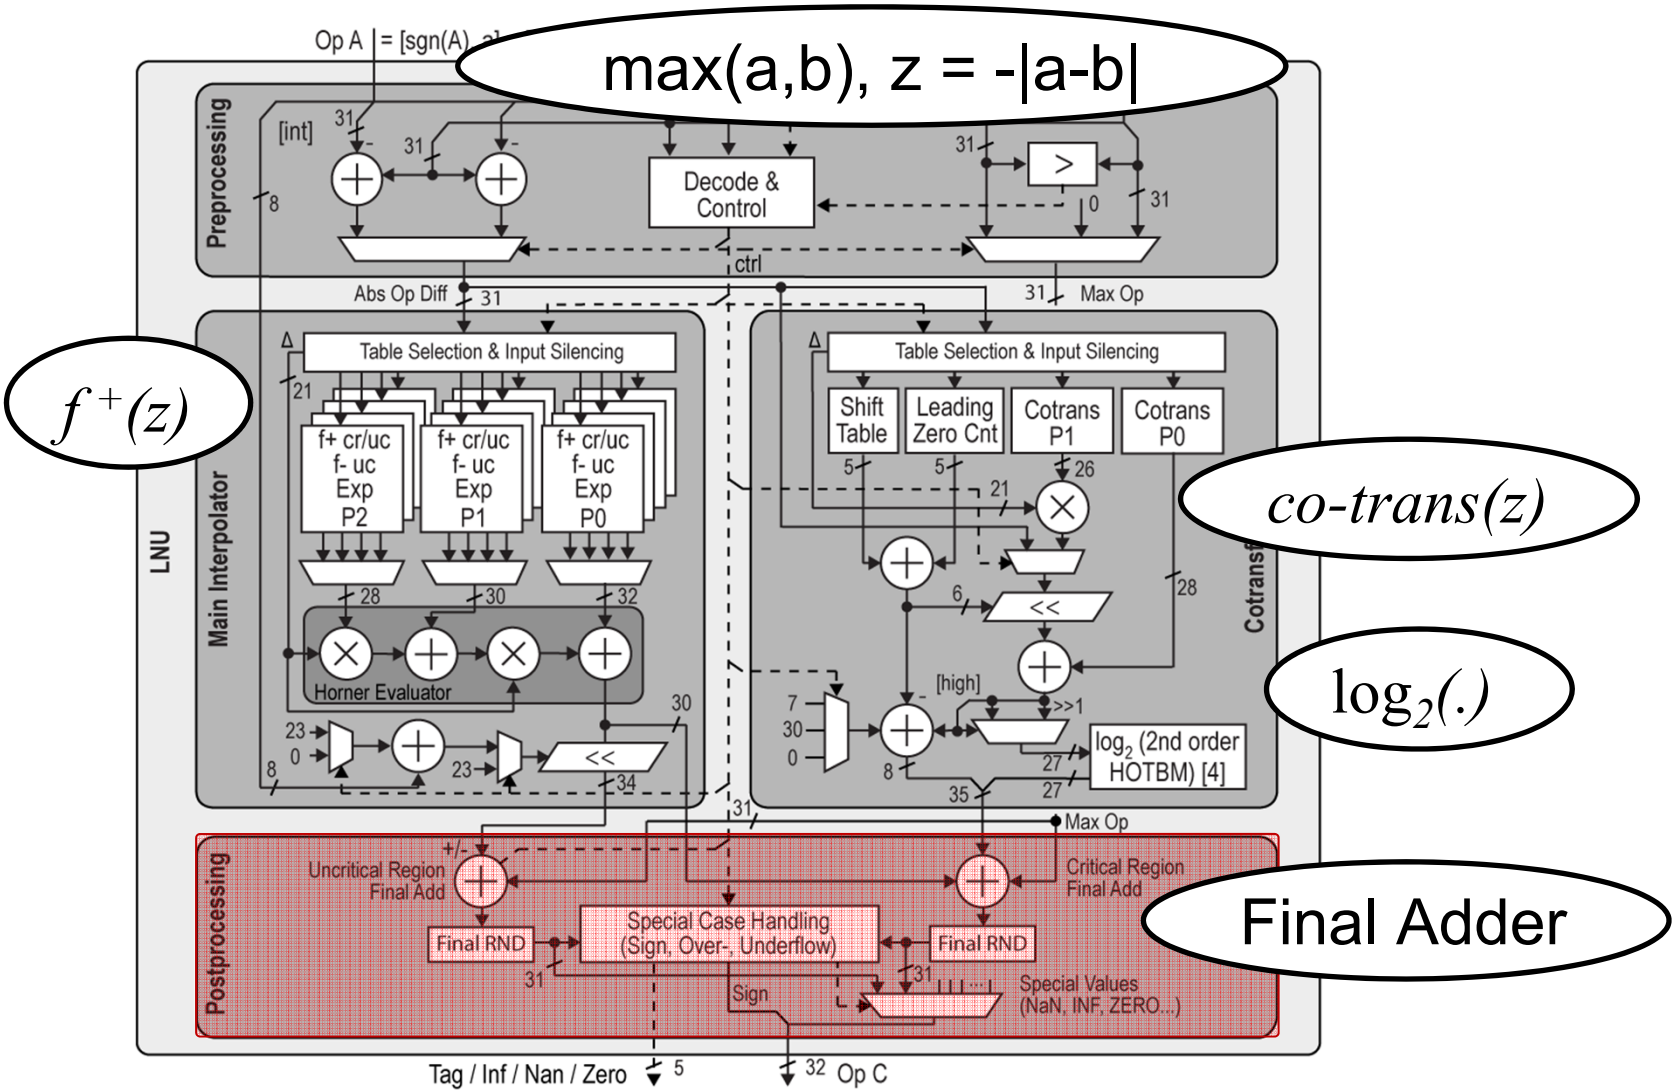

LNU Features & Comparison

	This Work	ELM [1,2]	ROM-less [2]
functionality	F2I, I2F, EXP, LOG ADD, SUB	ADD, SUB	ADD, SUB
max error [ulp]	0.479	0.454	0.498
LUT size [Kbit]	113.1	256.4	183.3
technology	65 nm	180 nm	180 nm
area [um^2]	57’264	904’943	589’357
Post-synthesis [kGE]	40	97	63
min delay [ns]	6	11.74	7.10
max delay [ns]	6	13.15	14.79




[1] J.N. Coleman et al. "The European Logarithmic Microprocessor" IEEE TC, 2008
[2] R.C. Ismail et al. "ROM-less LNS" IEEE ARITH, 2011

Hardware Platform

- Parallel Ultra-Low-Power (PULP) Platform [3]
  - 4x 32b OpenRISC Cores with 1kB private I$
  - 16 kByte L2 memory
  - 16 kByte shared L1 (TCDM)
  - SPI peripheral

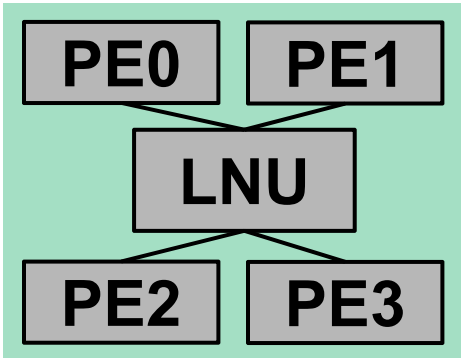

- Fabricated chip configurations:
  - **1 Shared LNU (shown)**
  - 4 Private FPUs (reference)

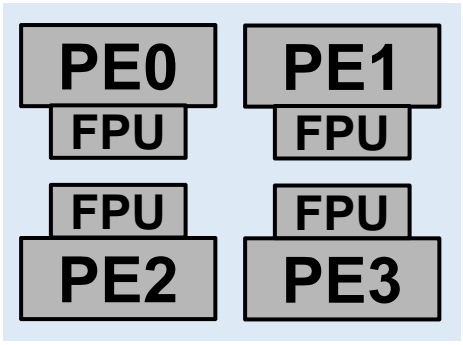

http://pulp.ethz.ch/

[3] D.Rossi et. al., "A -1.8V to 0.9V Body Bias, 60 GOPS/W 4-core Cluster in low-power 28nm UTBB FD-SOI technology," S3S, 2015"

PULP Architecture with shared LNU

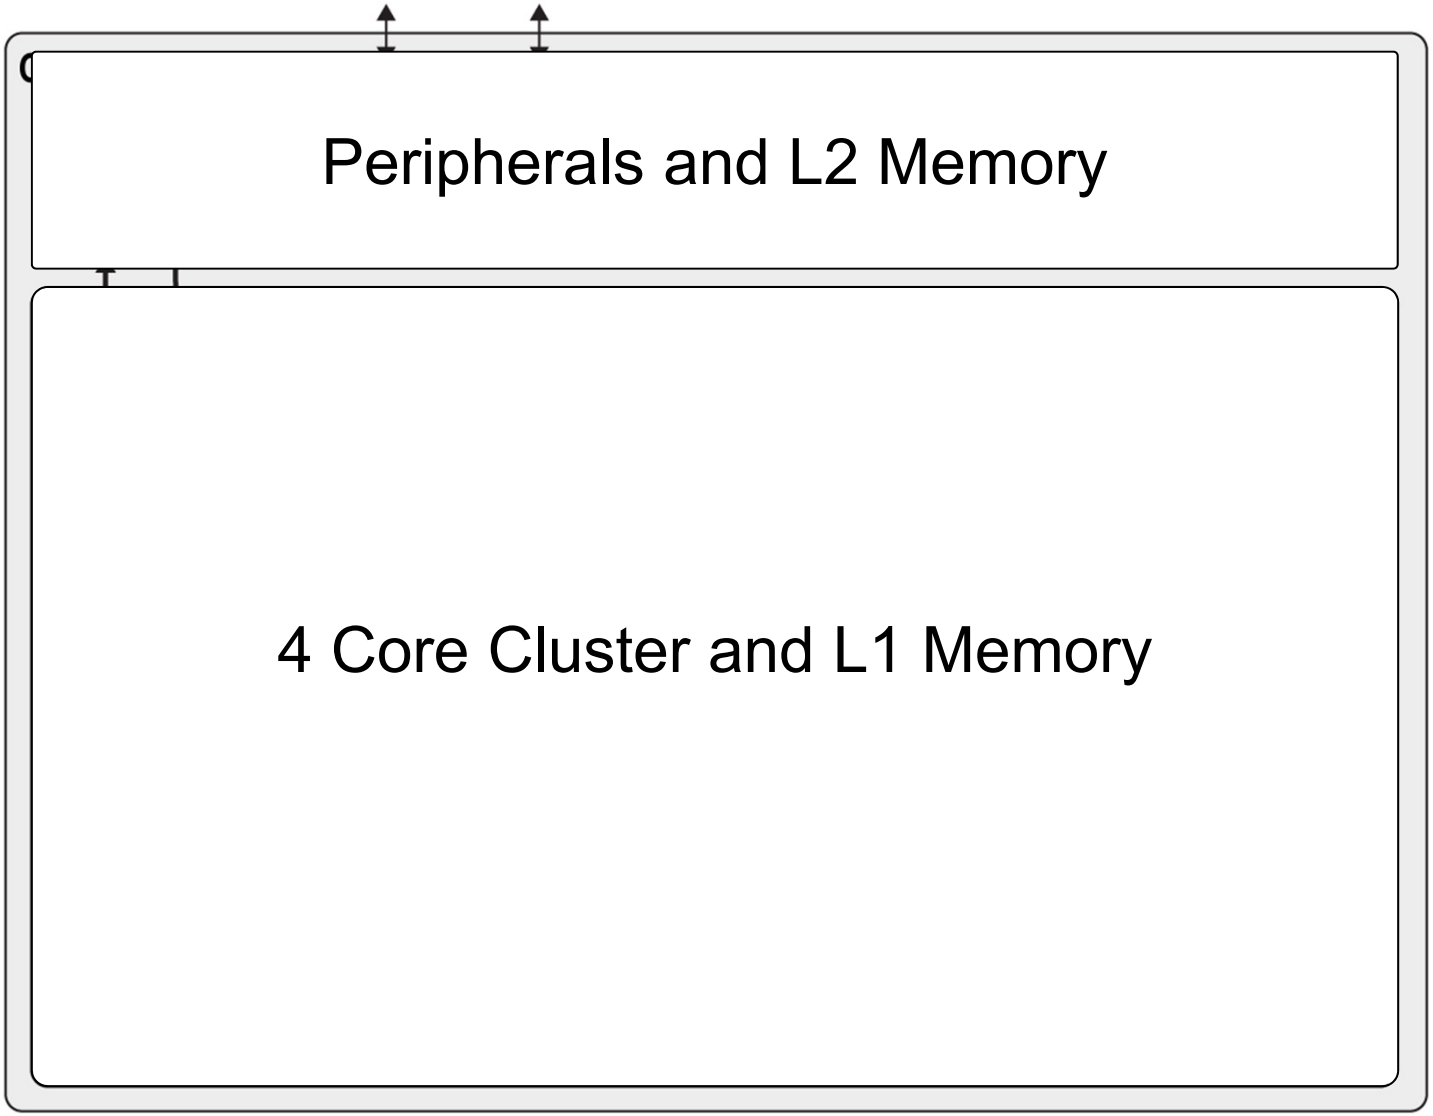

PULP Architecture with shared LNU

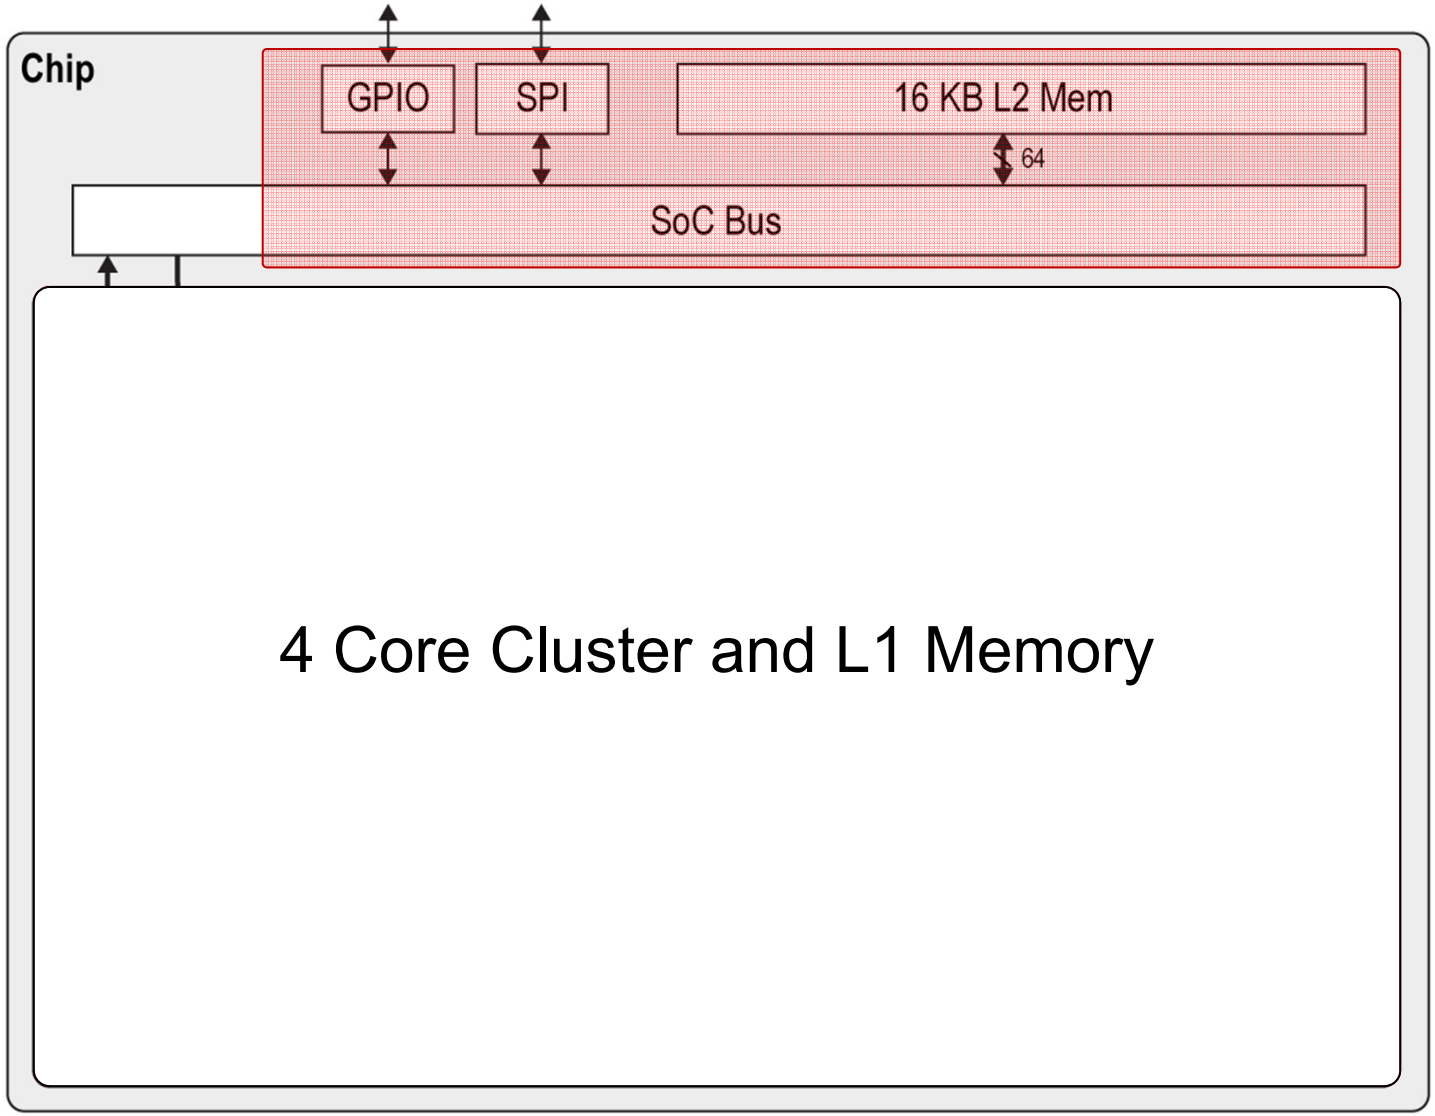

PULP Architecture with shared LNU

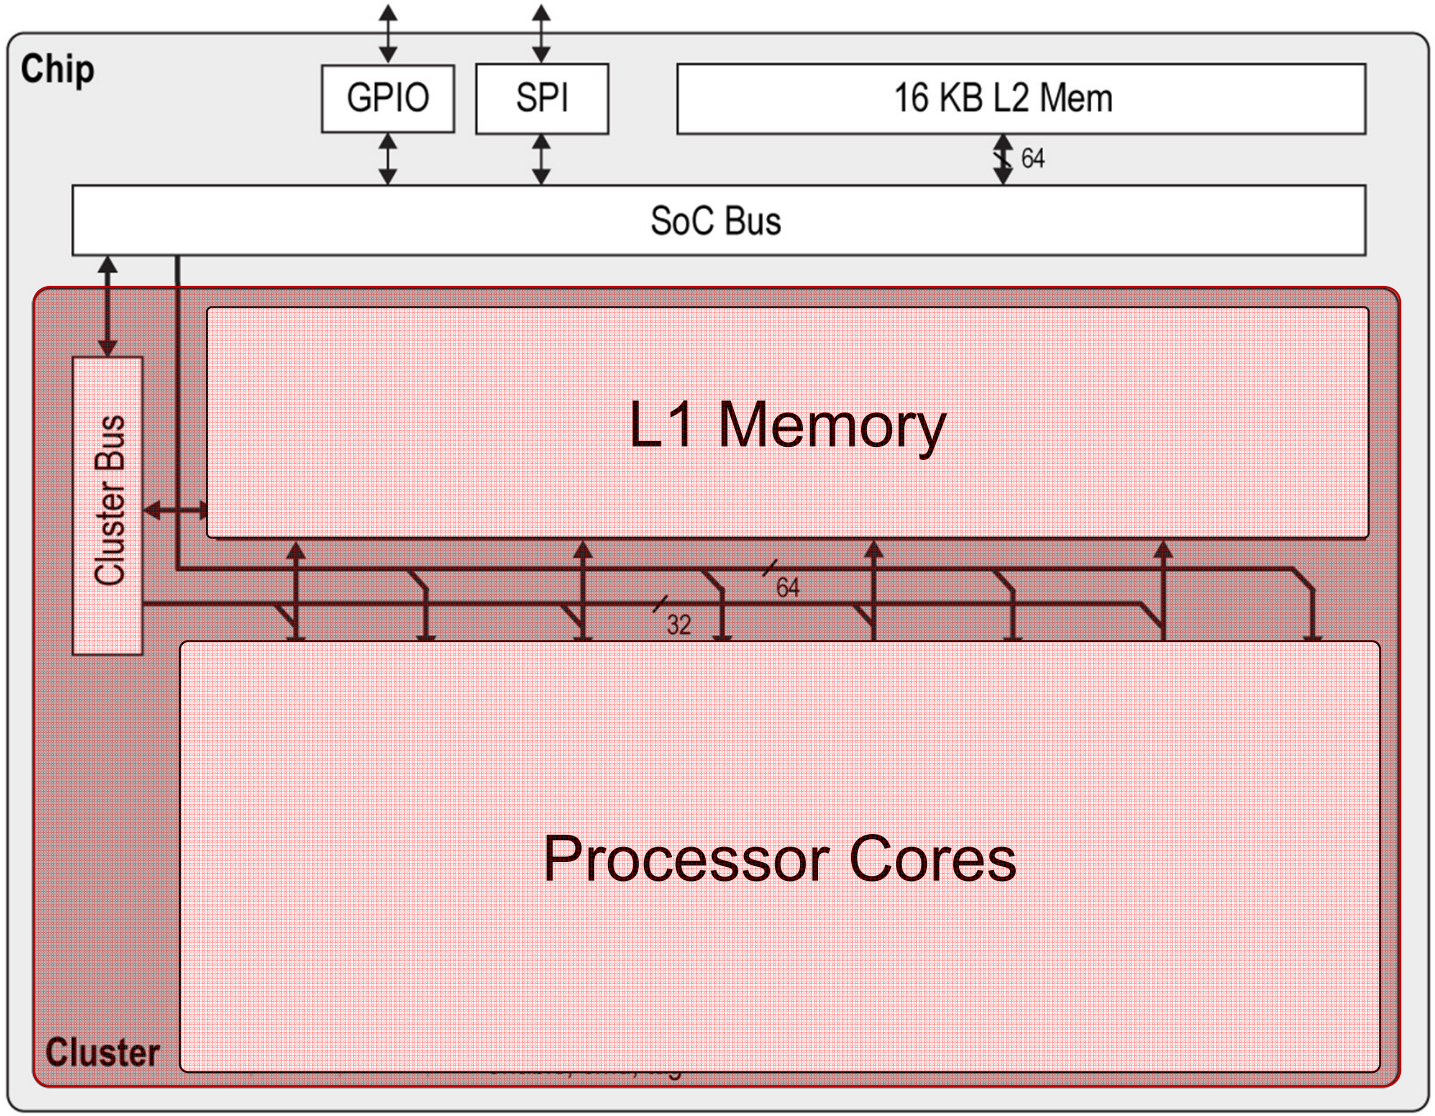

PULP Architecture with shared LNU

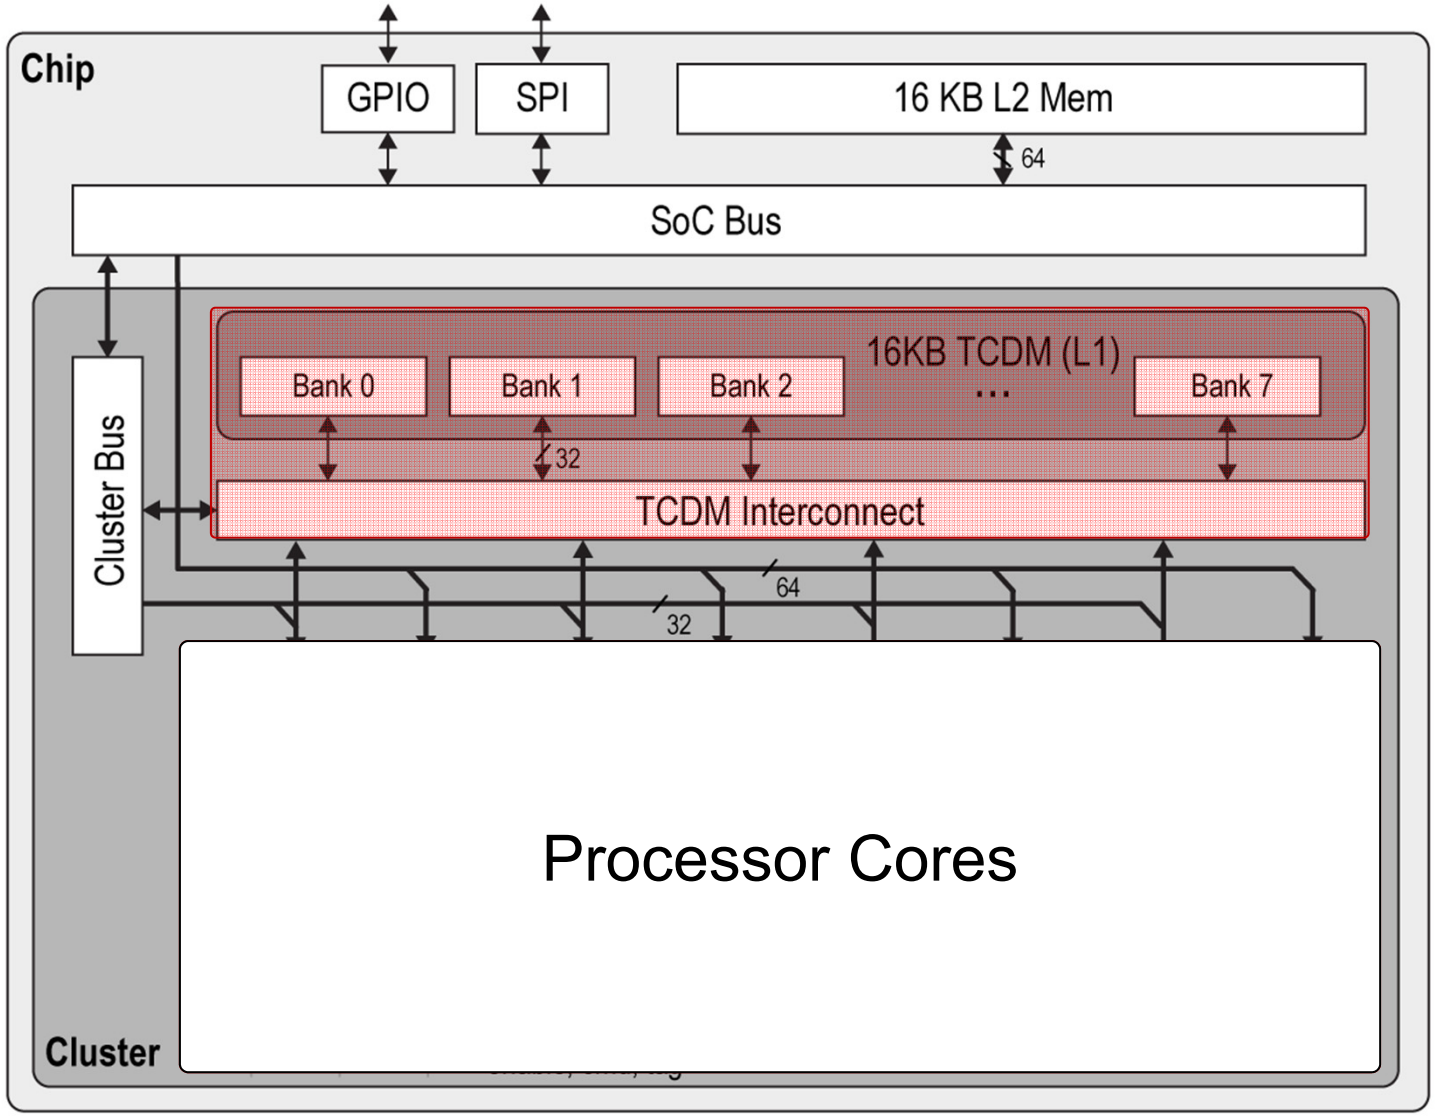

PULP Architecture with shared LNU

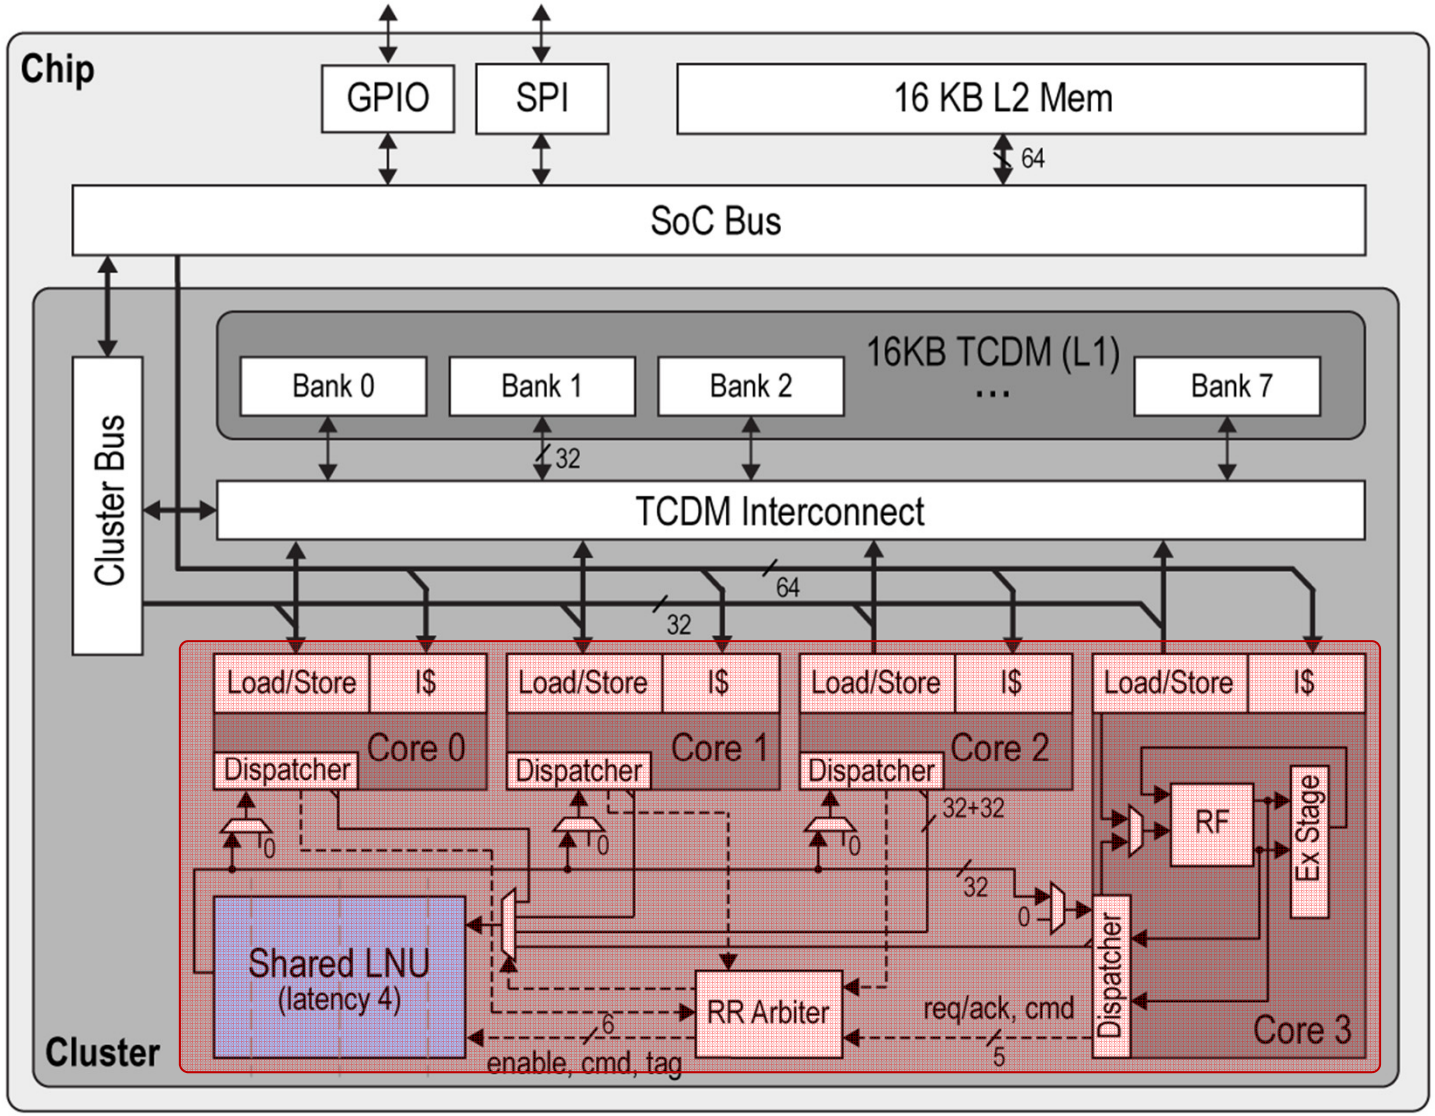

PULP Architecture with shared LNU

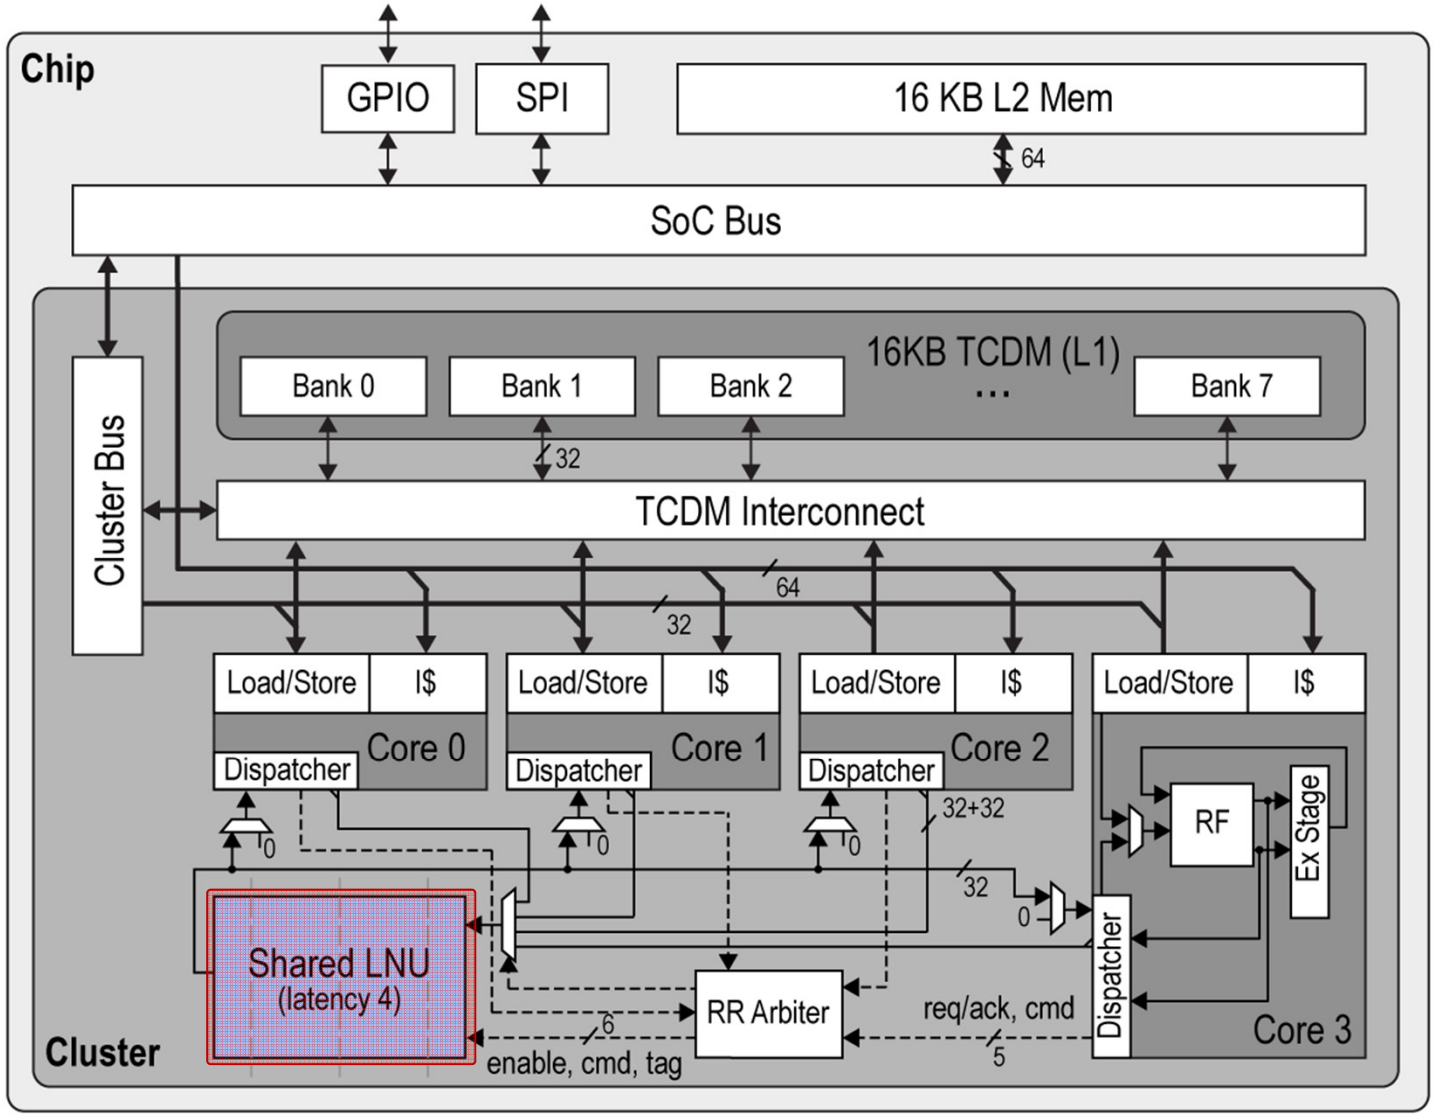

Energy Efficiency of FP-Instructions

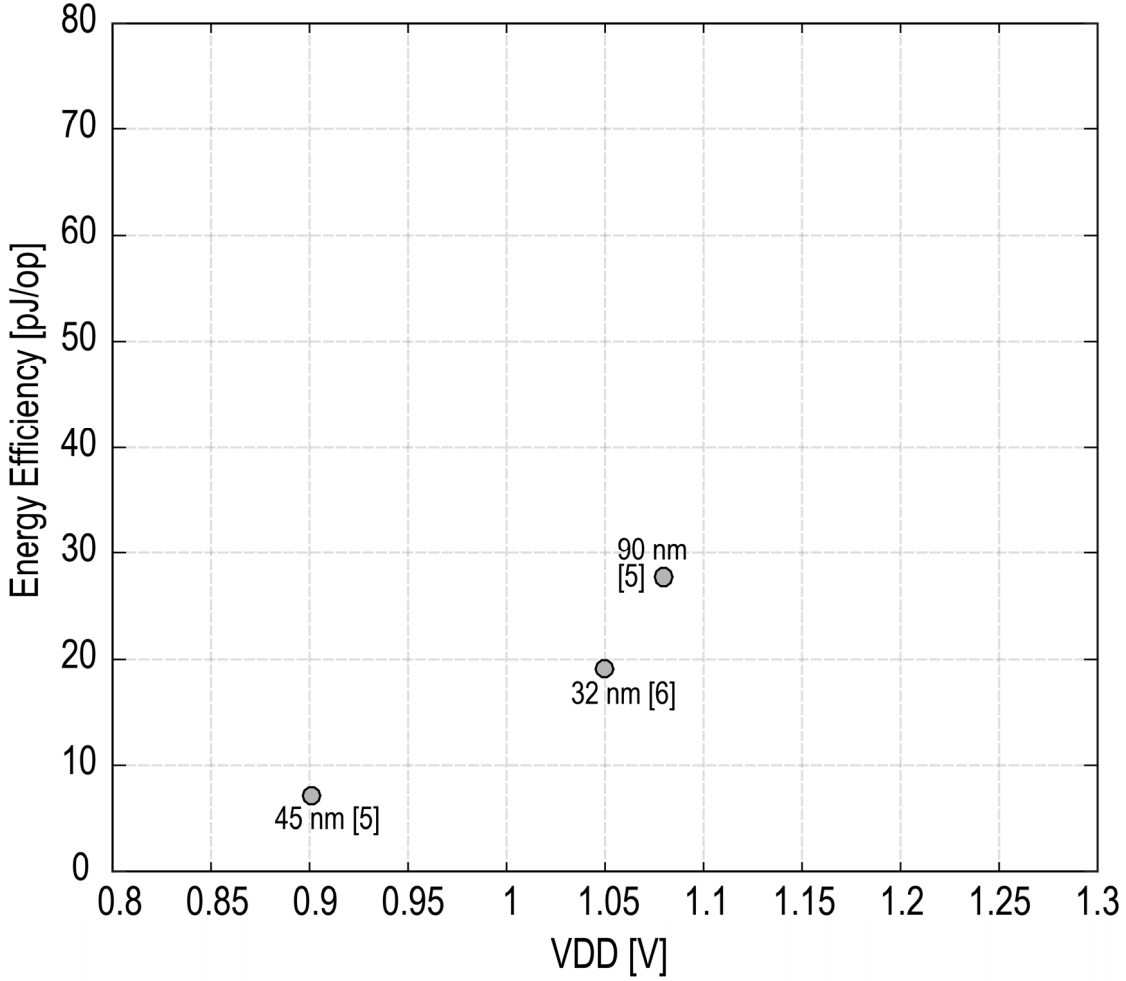

[5] S. Galal, M.Horowitz, "Energy-Efficient Floating-Point Unit Design," IEEE TC 2011

[6] H. Kaul et al., "A 1.45 GHz 52-to-162GFLOPS/W variable-precision floating point fused multiply-add unit with certainty tracking in 32nm CMOS." ISSCC'12

Measured Energy Efficiency of FP-Instructions

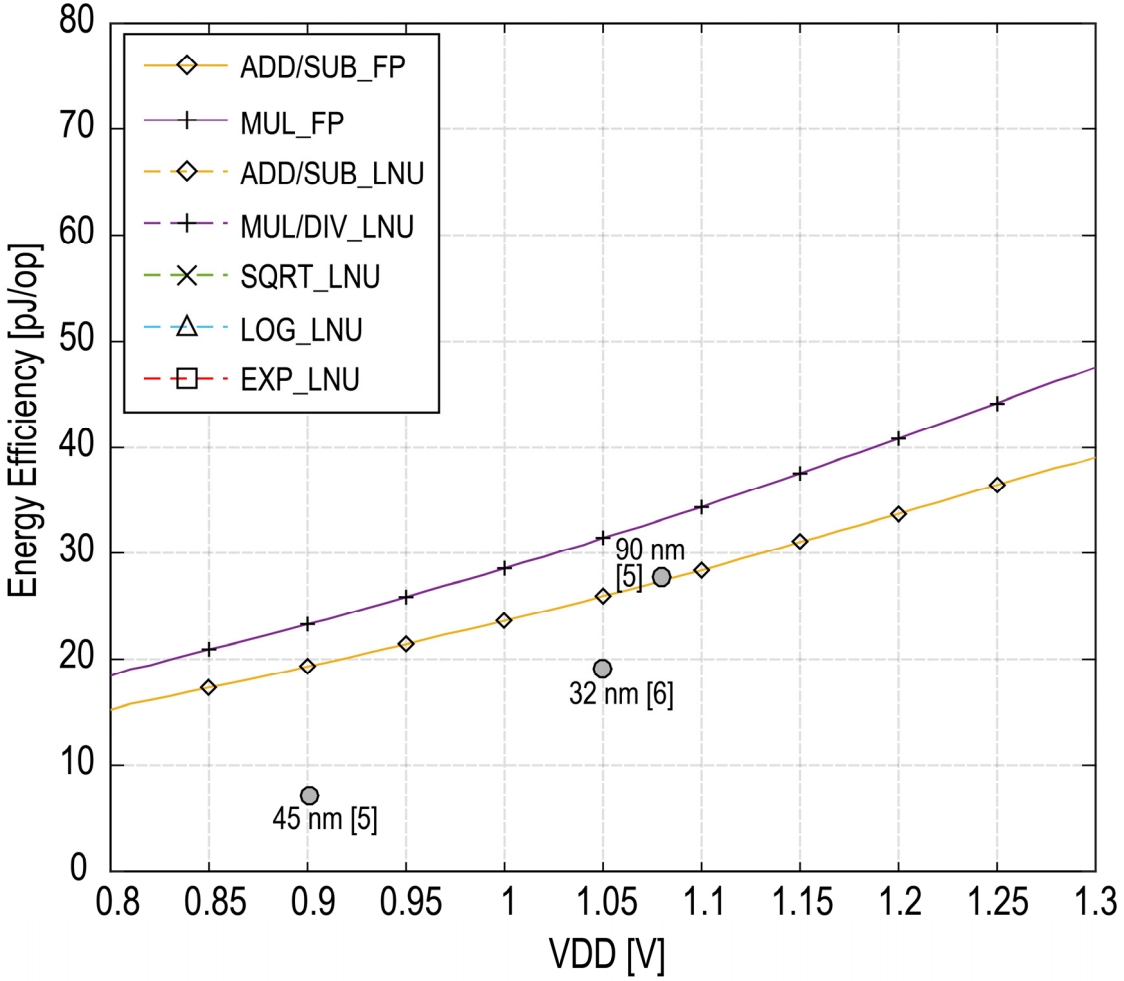

[5] S. Galal, M.Horowitz, "Energy-Efficient Floating-Point Unit Design," IEEE TC 2011

[6] H. Kaul et al., "A 1.45 GHz 52-to-162GFLOPS/W variable-precision floating point fused multiply-add unit with certainty tracking in 32nm CMOS." ISSCC'12

Measured Energy Efficiency of FP-Instructions

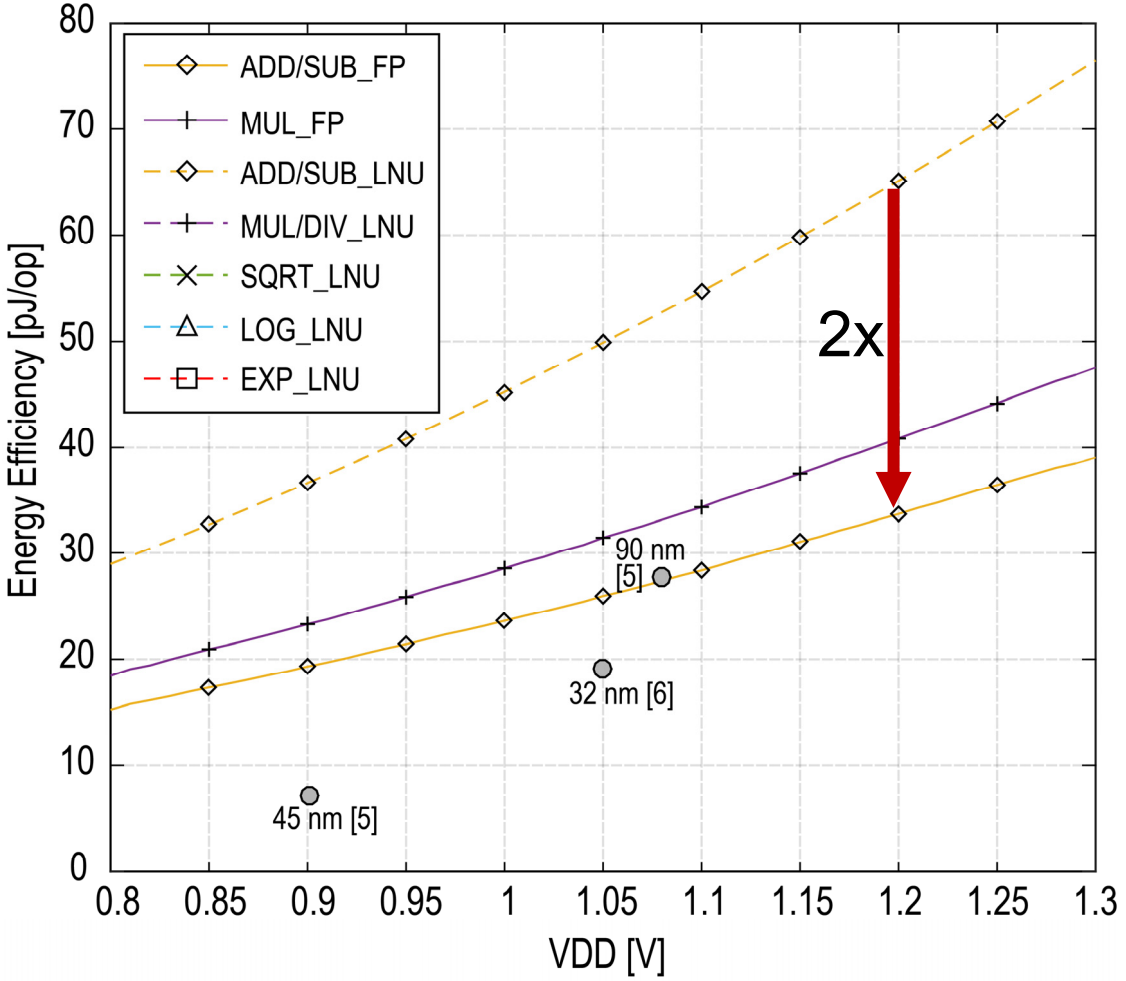

ADD/SUB: Expensive

[5] S. Galal, M.Horowitz, "Energy-Efficient Floating-Point Unit Design," IEEE TC 2011

[6] H. Kaul et al., "A 1.45 GHz 52-to-162GFLOPS/W variable-precision floating point fused multiply-add unit with certainty tracking in 32nm CMOS." ISSCC'12

Measured Energy Efficiency of FP-Instructions

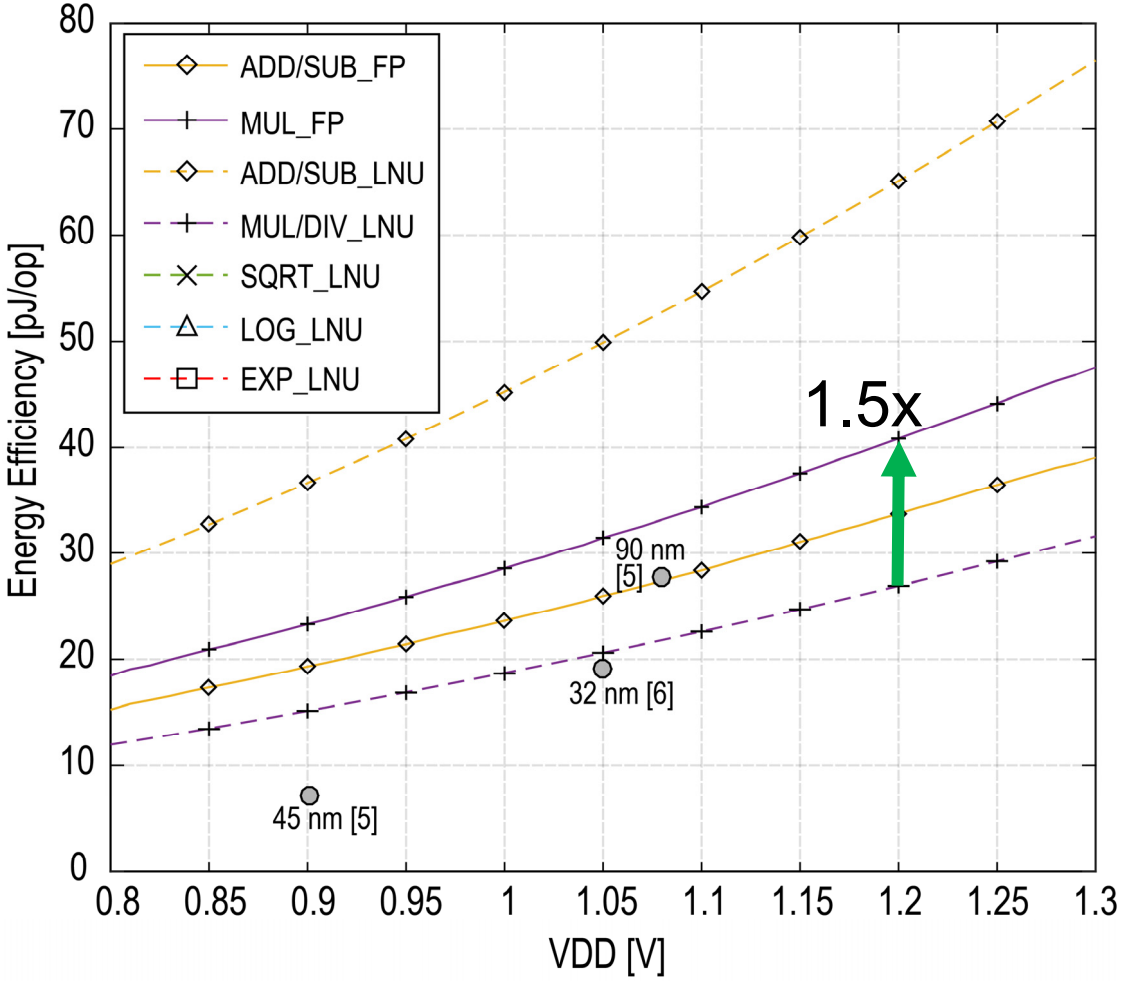

MUL/DIV: Cheap and single cycle!

[5] S. Galal, M.Horowitz, "Energy-Efficient Floating-Point Unit Design," IEEE TC 2011

[6] H. Kaul et al., "A 1.45 GHz 52-to-162GFLOPS/W variable-precision floating point fused multiply-add unit with certainty tracking in 32nm CMOS." ISSCC'12

Measured Energy Efficiency of FP-Instructions

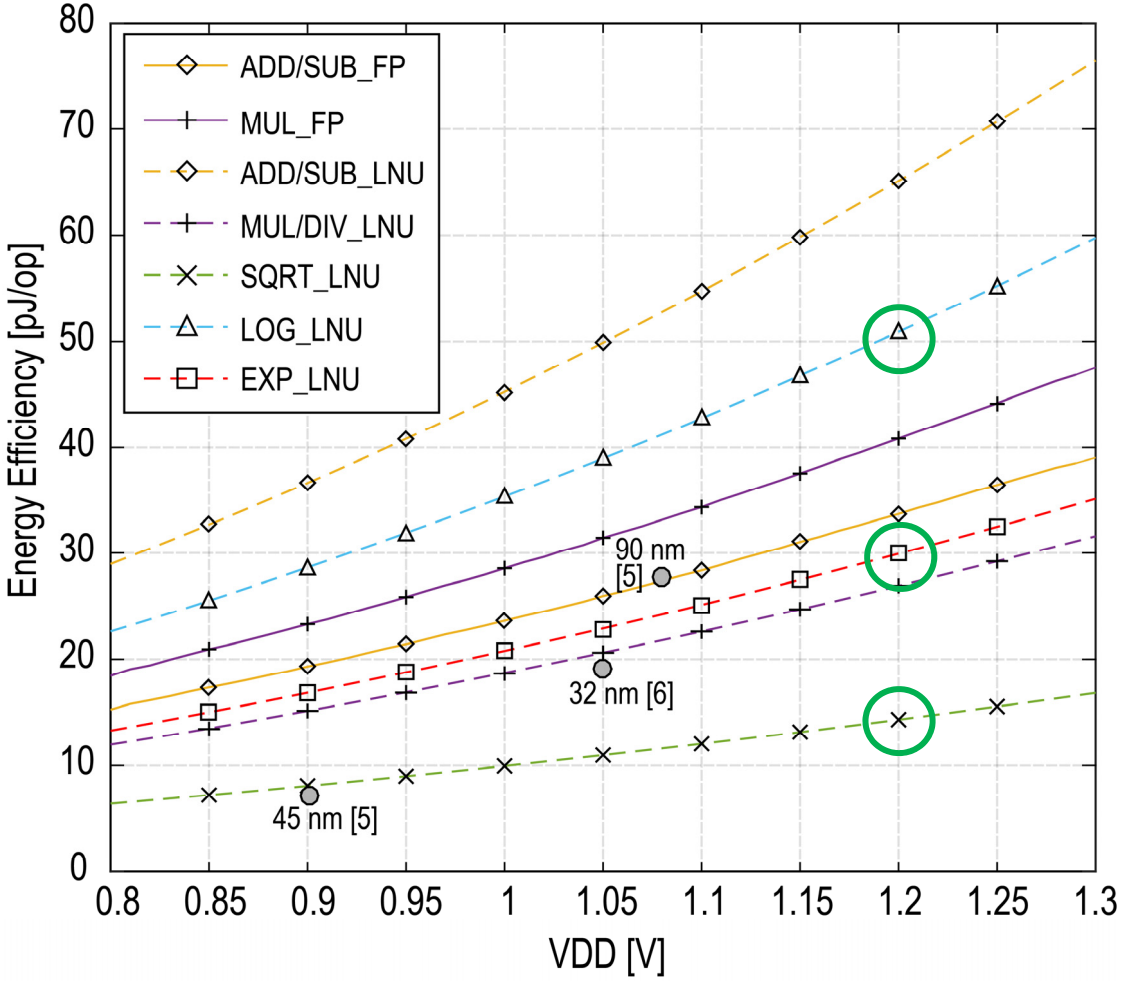

EXP/LOG: Single Instruction

SQRT: Cheap and single cycle!

[5] S. Galal, M.Horowitz, "Energy-Efficient Floating-Point Unit Design," IEEE TC 2011

[6] H. Kaul et al., "A 1.45 GHz 52-to-162GFLOPS/W variable-precision floating point fused multiply-add unit with certainty tracking in 32nm CMOS." ISSCC'12

Benchmark Kernels

Name	Details
AXPY	BLAS1 kernel, [100x1] vectors
GEMV	BLAS2 kernel, [10x10] matrices
GEMM	BLAS3 kernel, [10x10] matrices
Hom2D	2D projective transforms (stereo)
ProjErr2D	2D reprojection error of projective transforms (stereo)
Dist3D	3D Distance computations

Name	Details
Chol	Cholesky Decomposition, [10x10] matrices
DCT-II	32-point DCTs
QR	QR Decomposition, [10x10] matrices
LN	Element-wise Natural Logarithm
RBF	2D regression function with 25 Gaussian kernels

Benchmark Kernels

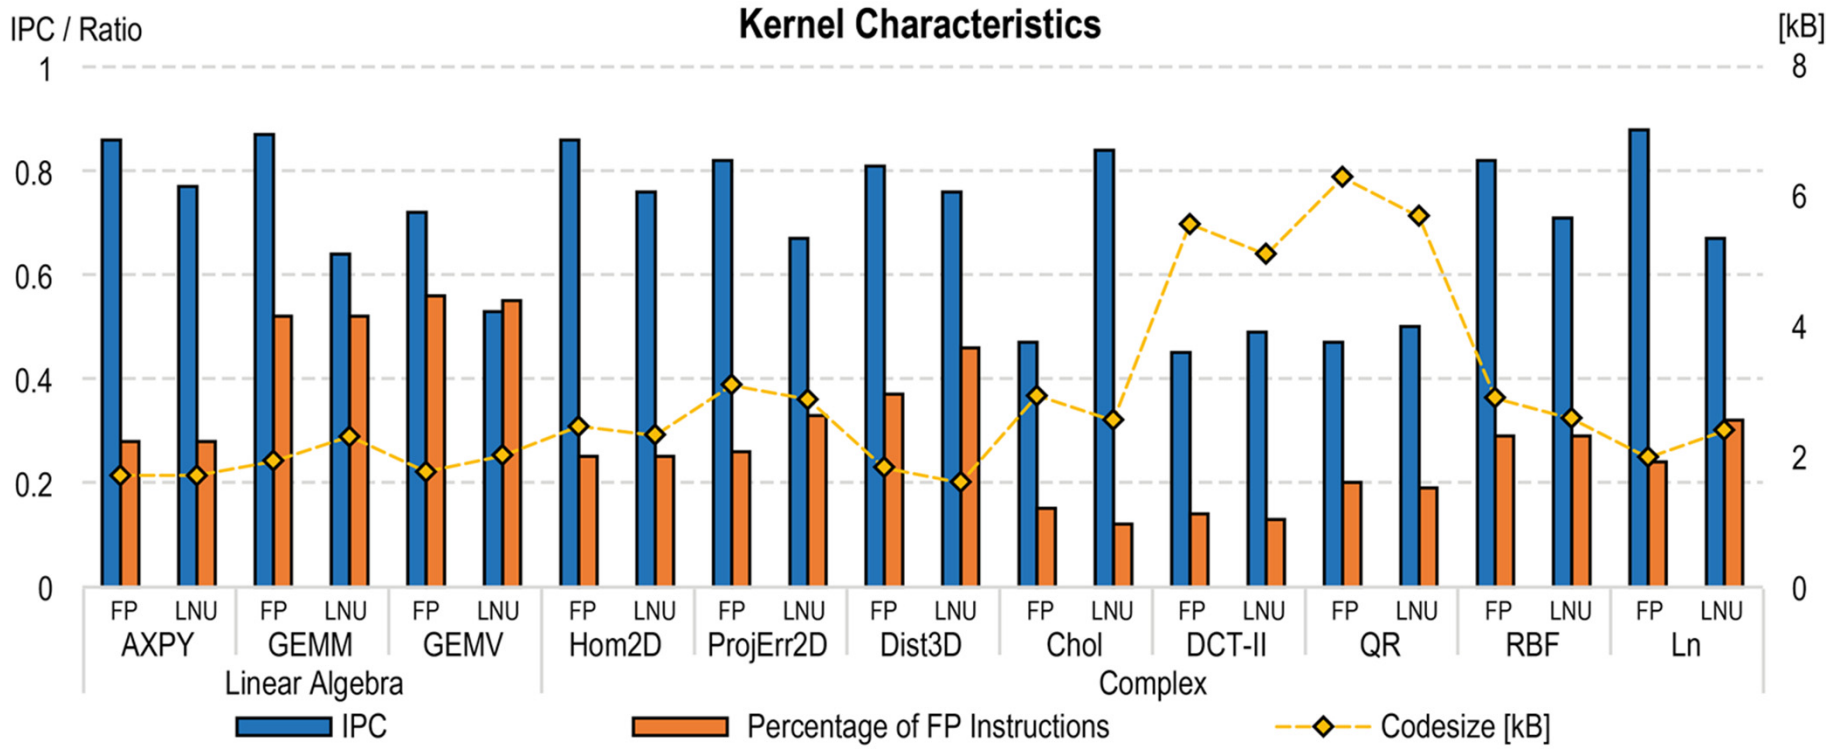

Energy Efficiency of Kernels (1)

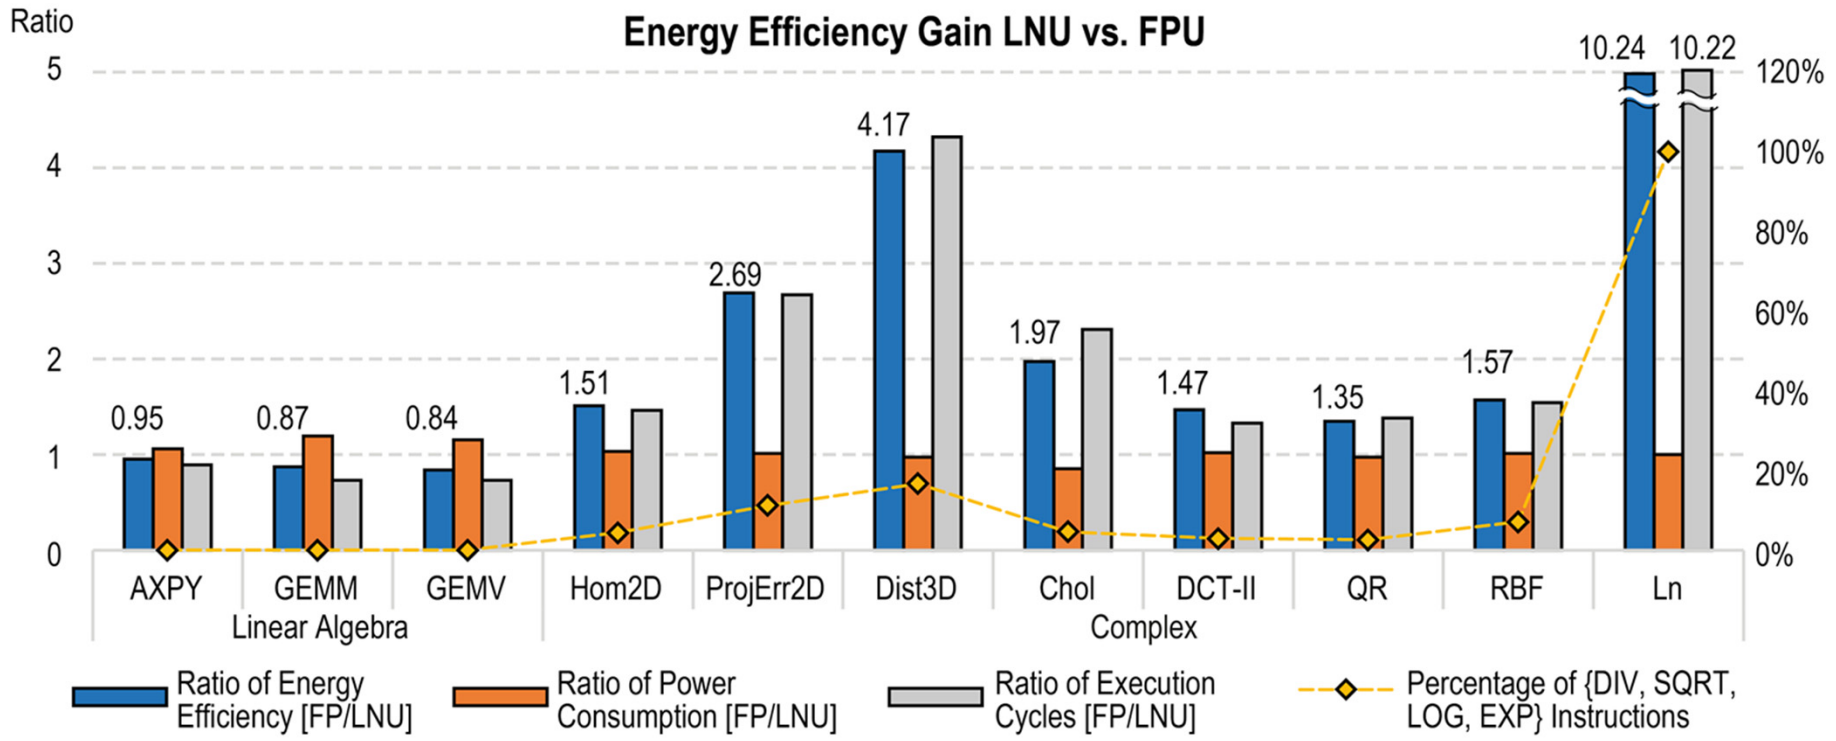

Energy Efficiency of Kernels (2)

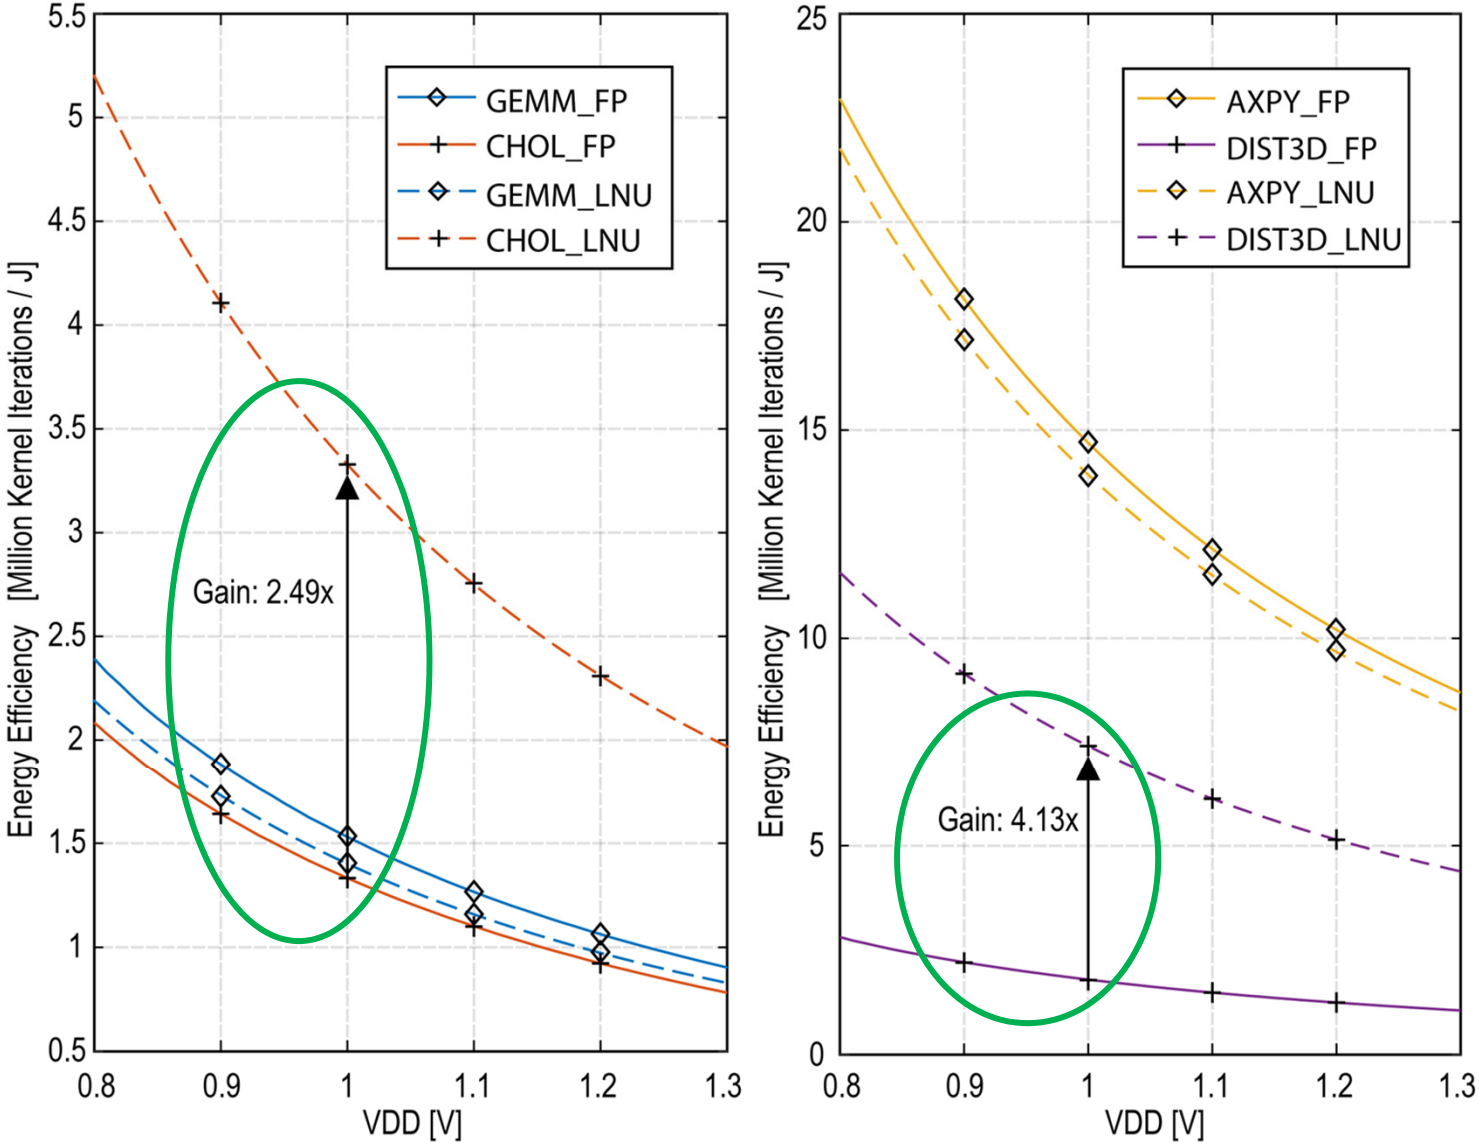

Chip Micrographs

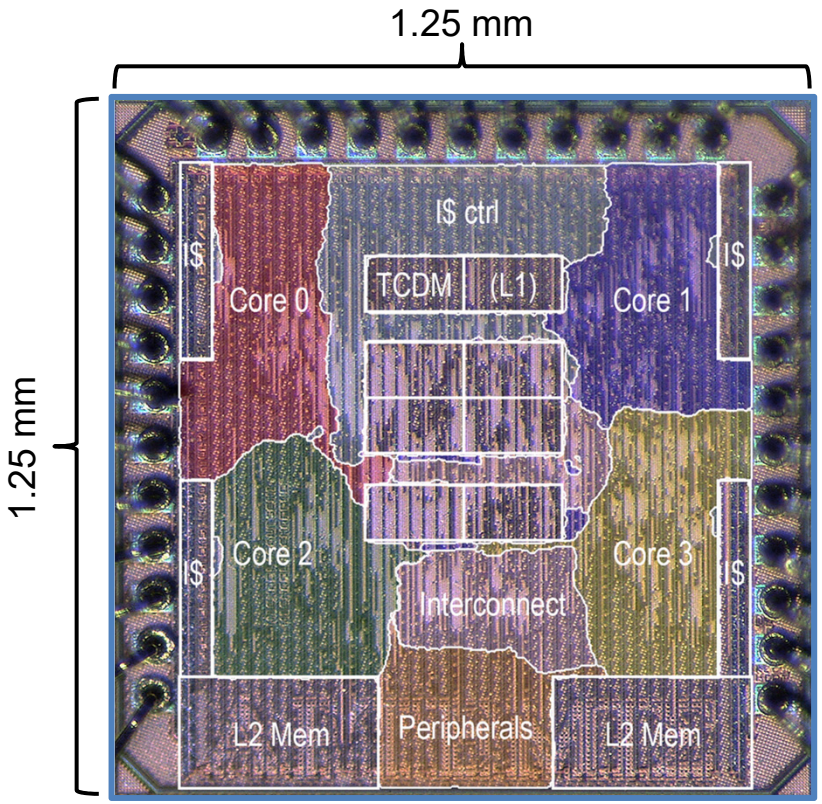

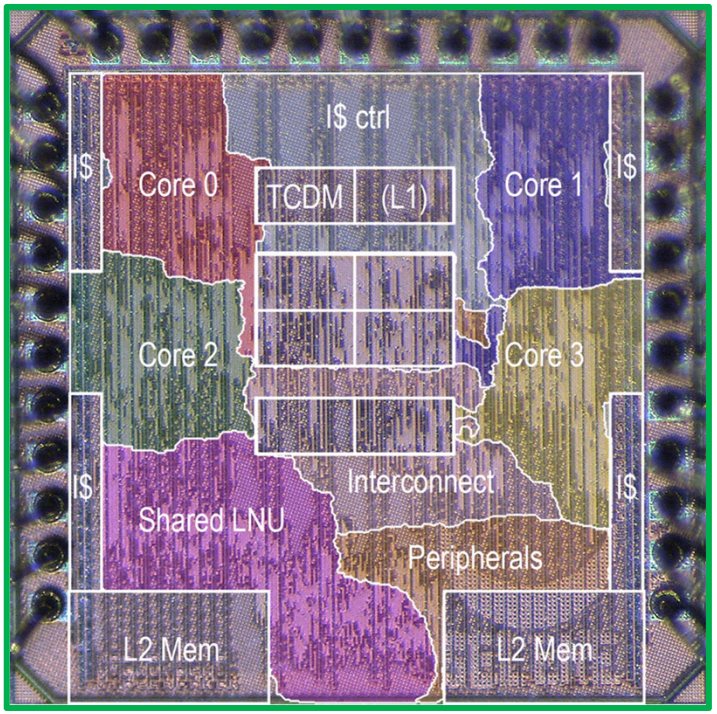

Area breakdown	Private FPU	Shared LNU
Total area [kGE]	719	749
FPU/LNU area [kGE]	4 x 10.8	53.1

Chip Datasheet & Comparison

Implementation Details	Private FPU	Shared LNU	ELM [1,2]
Technology	65nm LVT	65nm LVT	180nm
Max speed [MHz]	374	337	125
Functionality	+, -, *, casts	+, -, *, /, **exp**, **log**, casts	+, -, *
Precision (max err) [ulp]	0.5	0.478	0.454
Avg. LNU/FPU utilization	0.21	0.37	-
Avg. # stalls due to contentions	-	4%	-
Power @100MHz, 1.2V, 25°C [mW]	41.84	44.0	-
Leakage @ 1.2V, 25°C [mW]	2.823	3.019	-
Energy efficiency gain (complex)	1	**1.35 - 4.17**	

[1] J.N. Coleman et al. "The European Logarithmic Microprocessor" IEEE TC, 2008

[2] R.C. Ismail et al. "ROM-less LNS" IEEE ARITH, 2011

Conclusions

- **Shared** setting attractive for LNU
- **Fast SQRT, DIV, EXP, LOG** function evaluation!
- Up to **4.2x** more energy efficient than private FPU

Ongoing Work

- Interpolator improvements
- Vectorization
- Trigonometric Extensions

Q&A

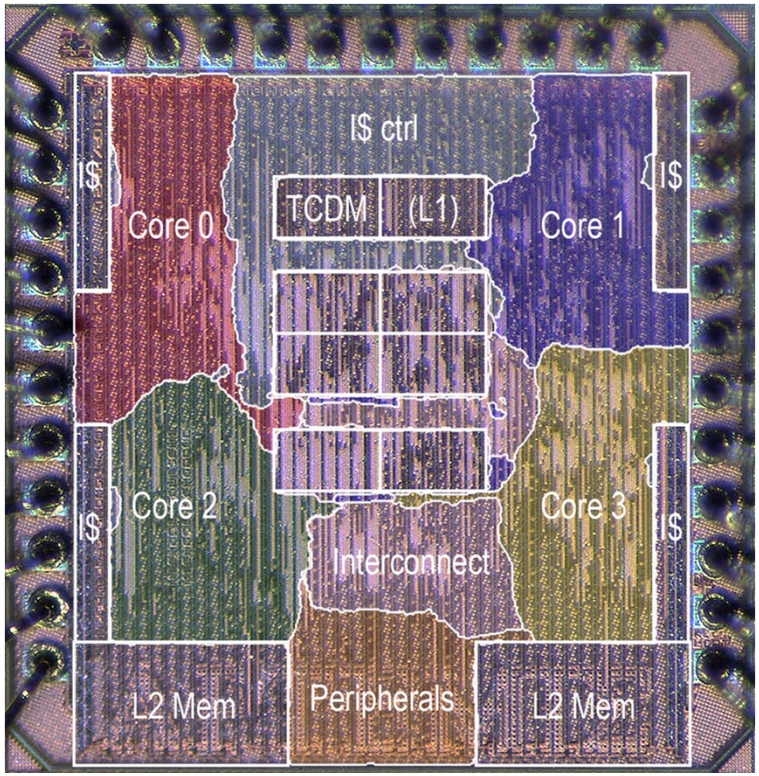

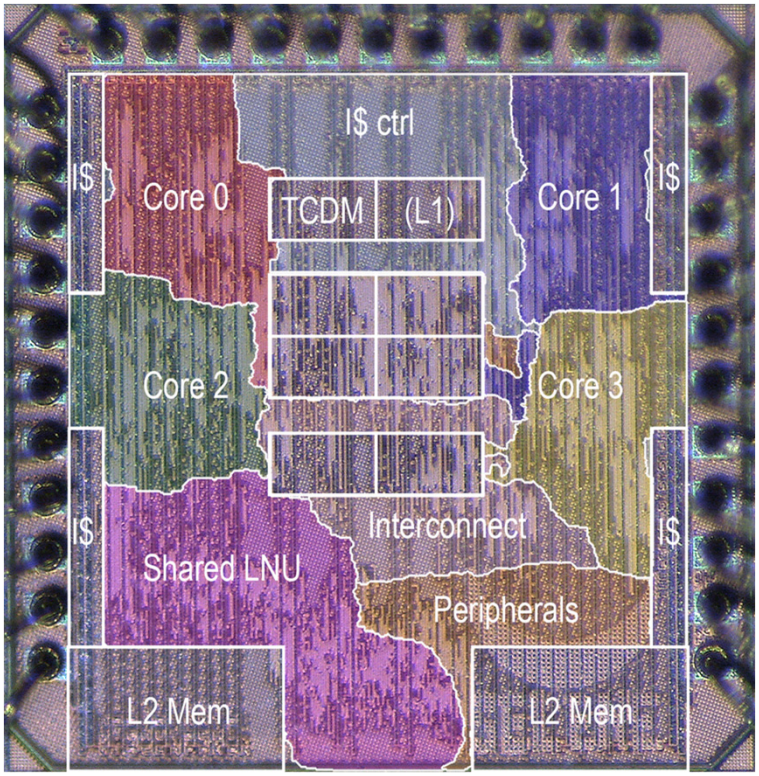

Acknowledgments:

- Nano Tera IcySoC project
- M. Burger, T. Gautschi, L. Müller, Y. Popoff, F. Scheidegger, F. Schuiki

A 65nm ReRAM-Enabled Nonvolatile Processor with 6× Reduction in Restore Time and 4× Higher Clock Frequency Using Adaptive Data Retention and Self-Write-Termination Nonvolatile Logic

Yongpan Liu[1], Zhibo Wang[1], Albert Lee[2,3], Fang Su[1], Chieh-Pu Lo[2], Zhe Yuan[1], Chien-Chen Lin[2], Qi Wei[1], Yu Wang[1], Ya-Chin King[2], Chrong-Jung Lin[2], Pedram Khalili[3], Kang-Lung Wang[3], Meng-Fan Chang[2], Huazhong Yang[1]

[1]Tsinghua University, Beijing,

[2]National Tsing Hua University, Hsinchu,

[3]University of California, Los Angeles

Outline

- **Background**
- Motivation
- System Architecture
  - 1-Macro Structure
  - Time-Domain Adaption
  - Space-Domain Adaption
- Measurement & Analysis

Power Challenge for IoTs

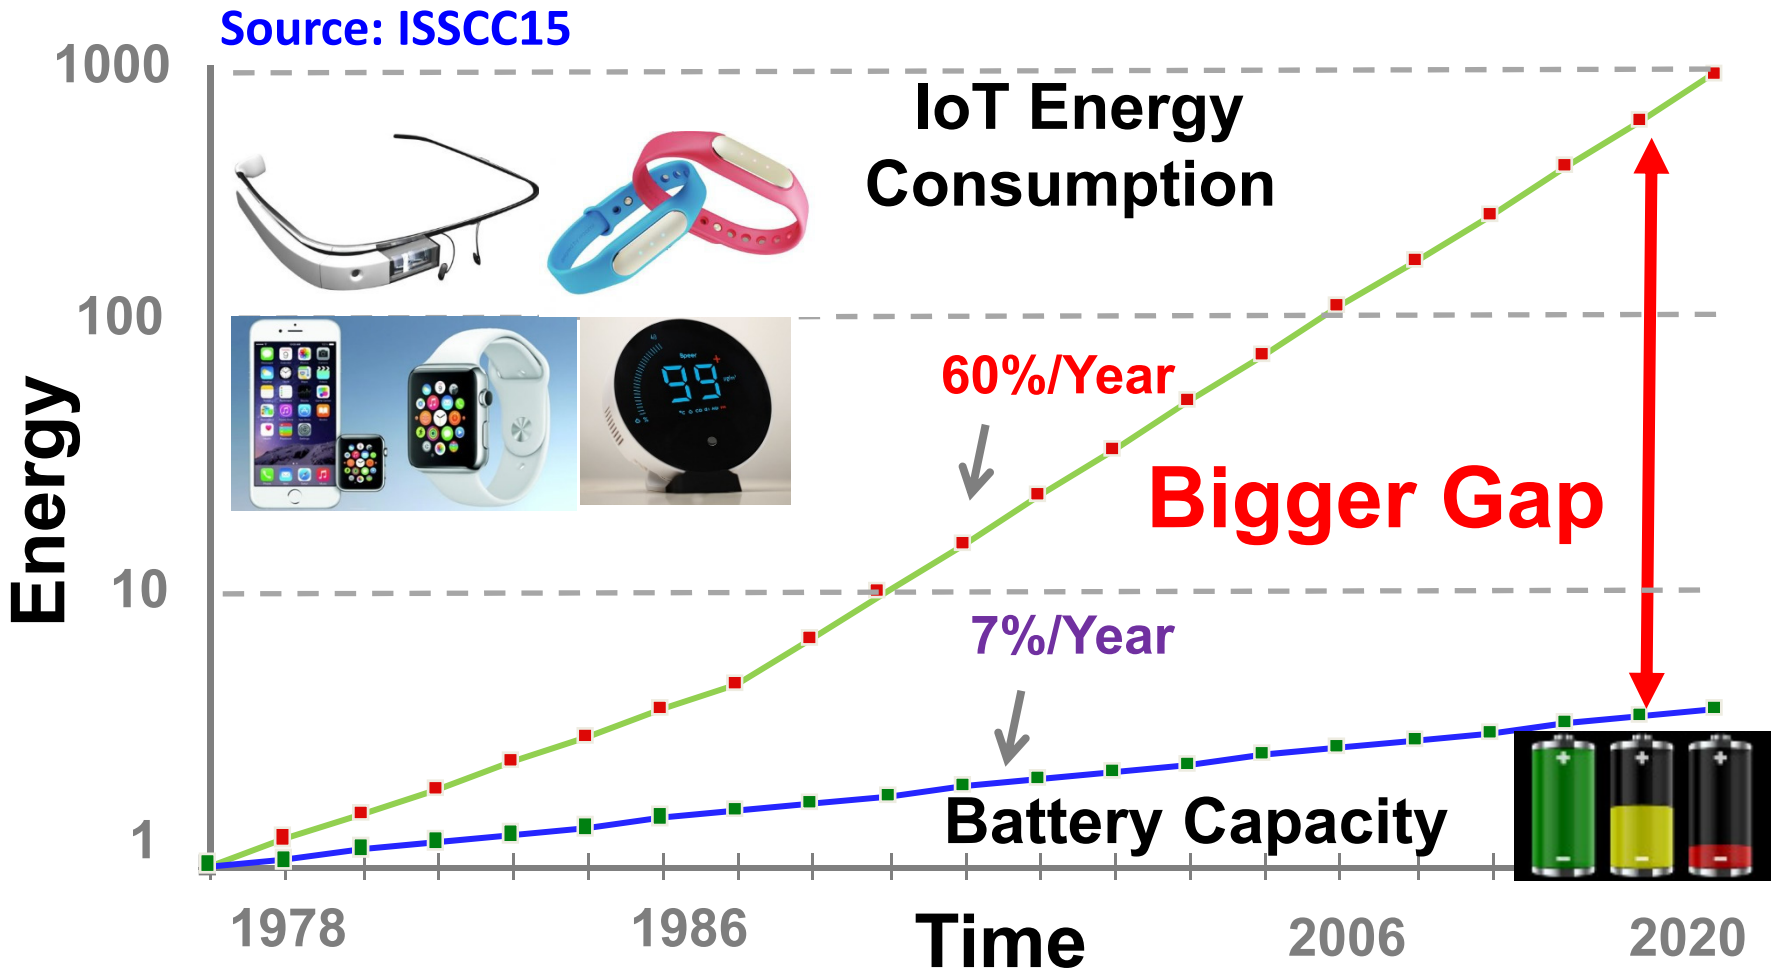

Can We Survive Without Charging Battery?

Battery-Less Energy Harvesting System

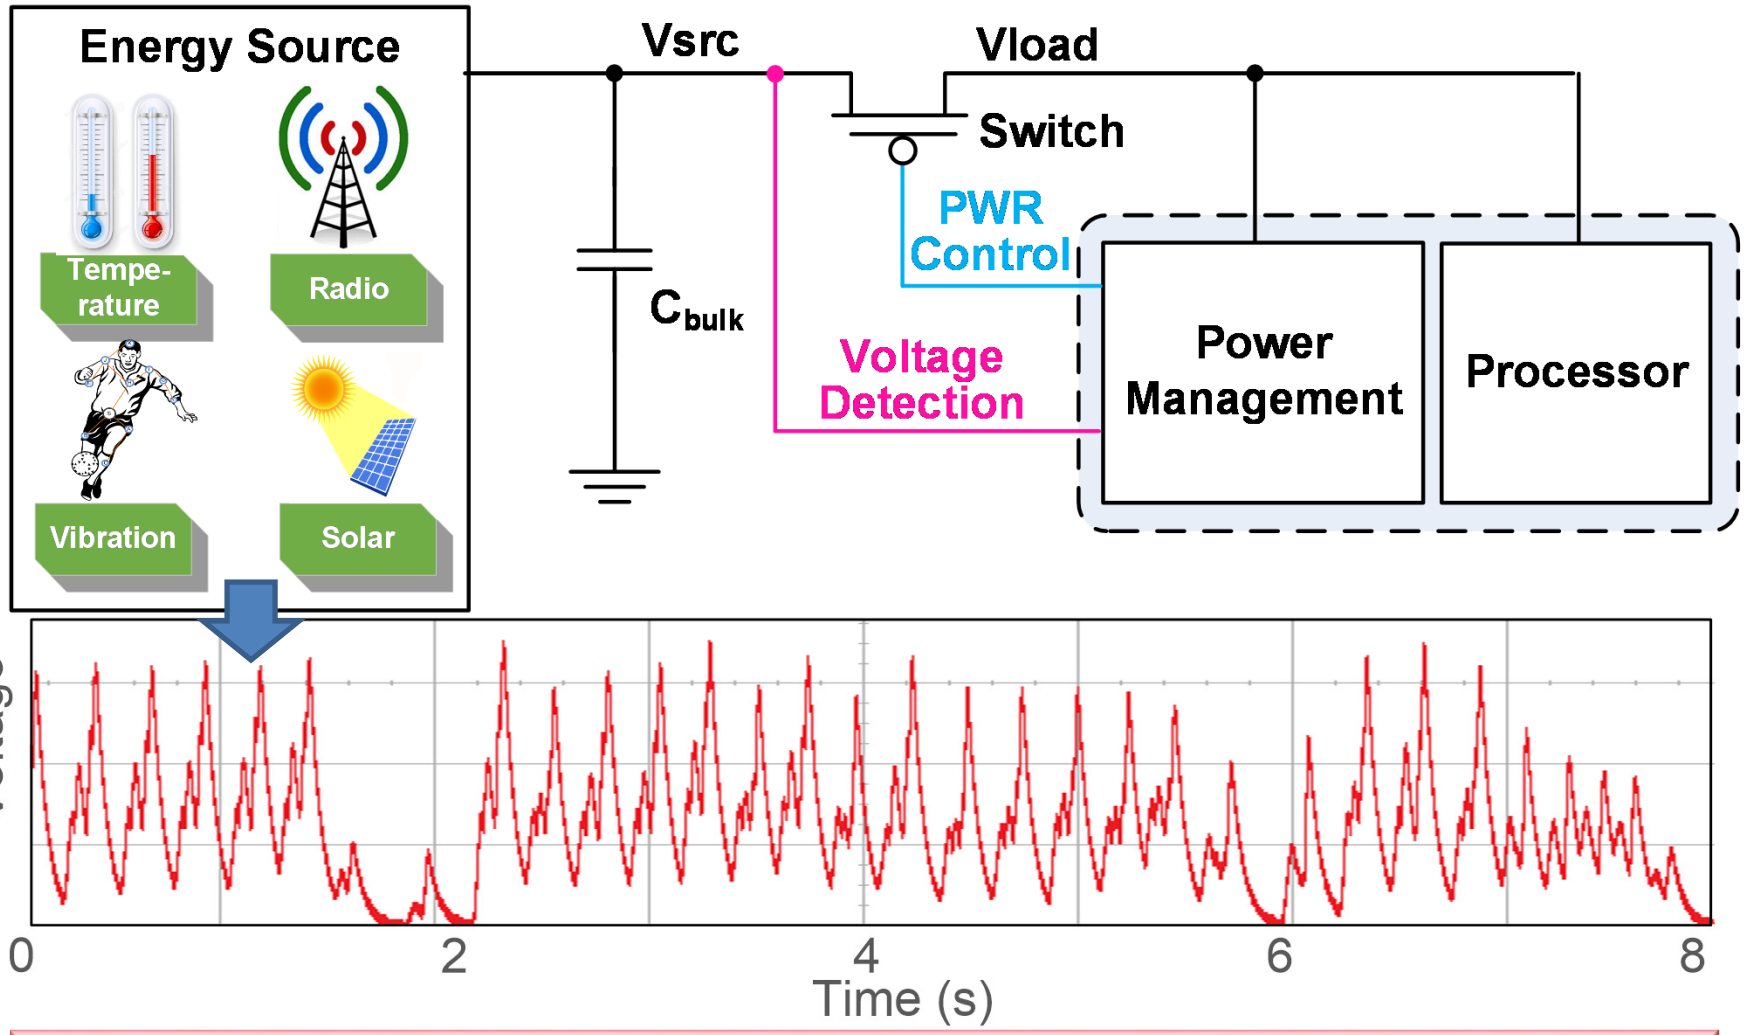

Energy Harvesting Is Promising But Unstable!

Computation under Energy Harvesting

Source: HPCA15 Best Paper Award

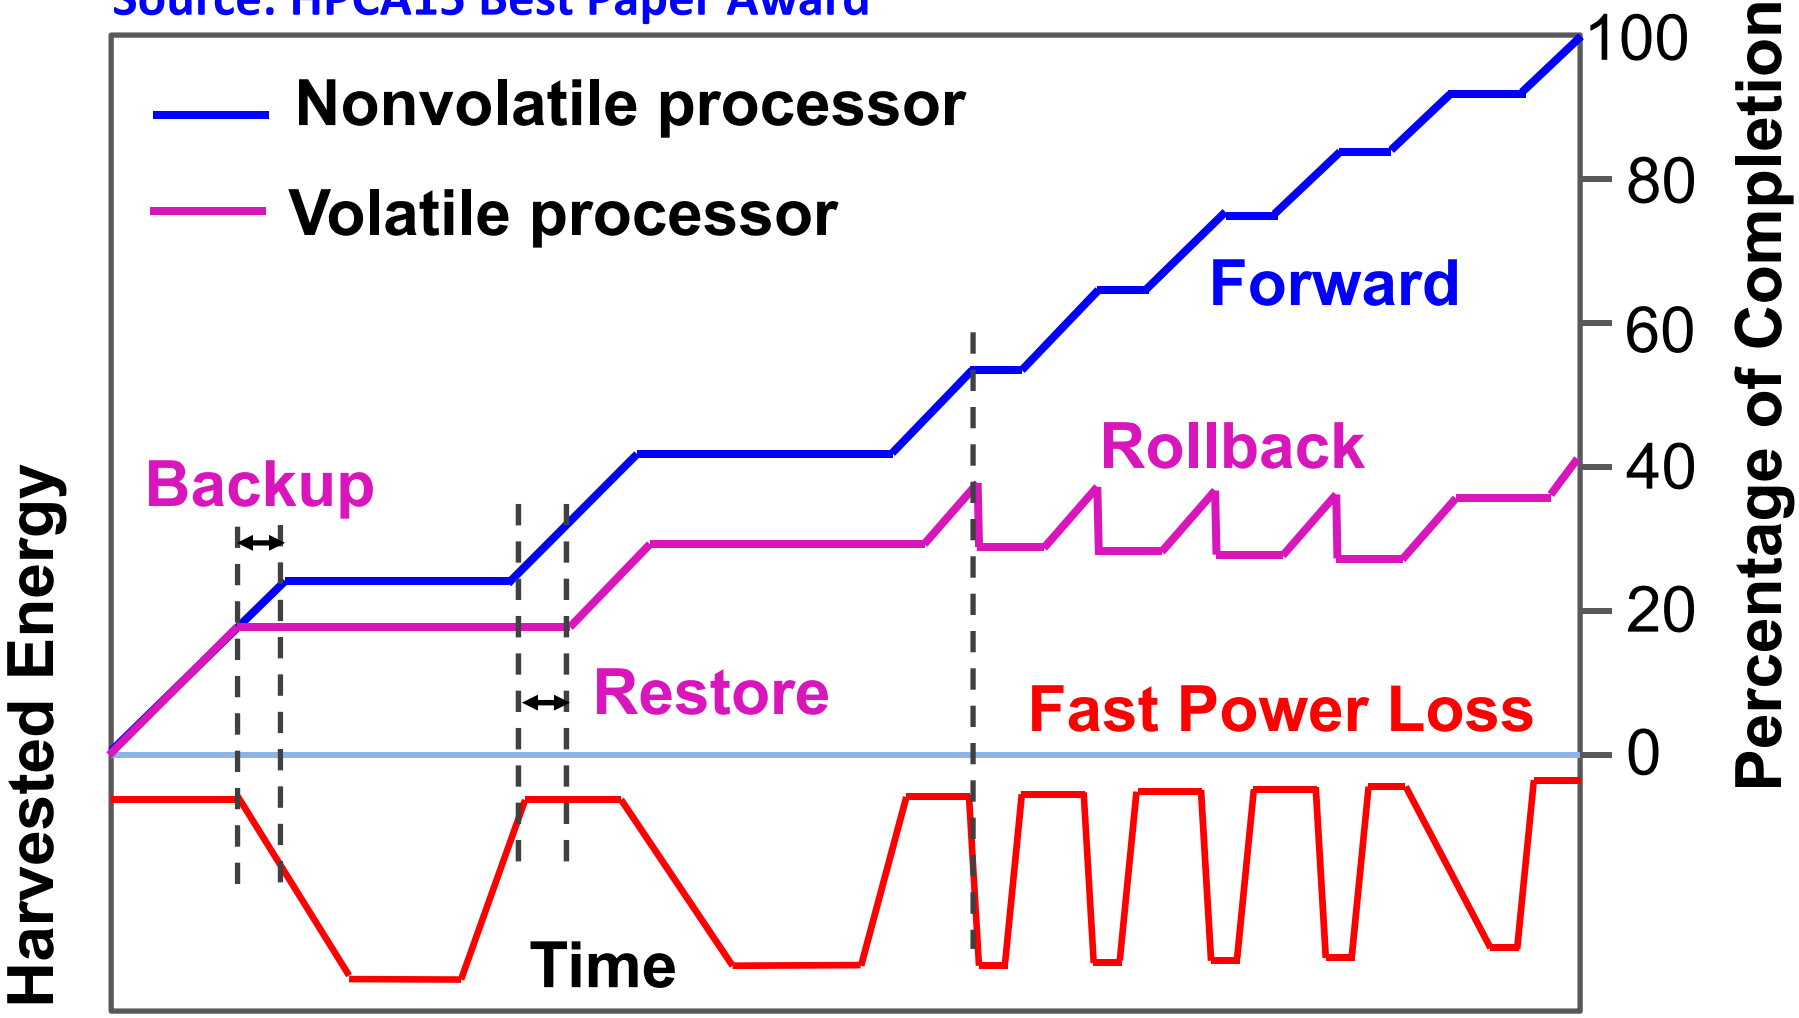

Processors Suffer from Frequent Rollback!

Nonvolatile Processor (NVP) for IoTs

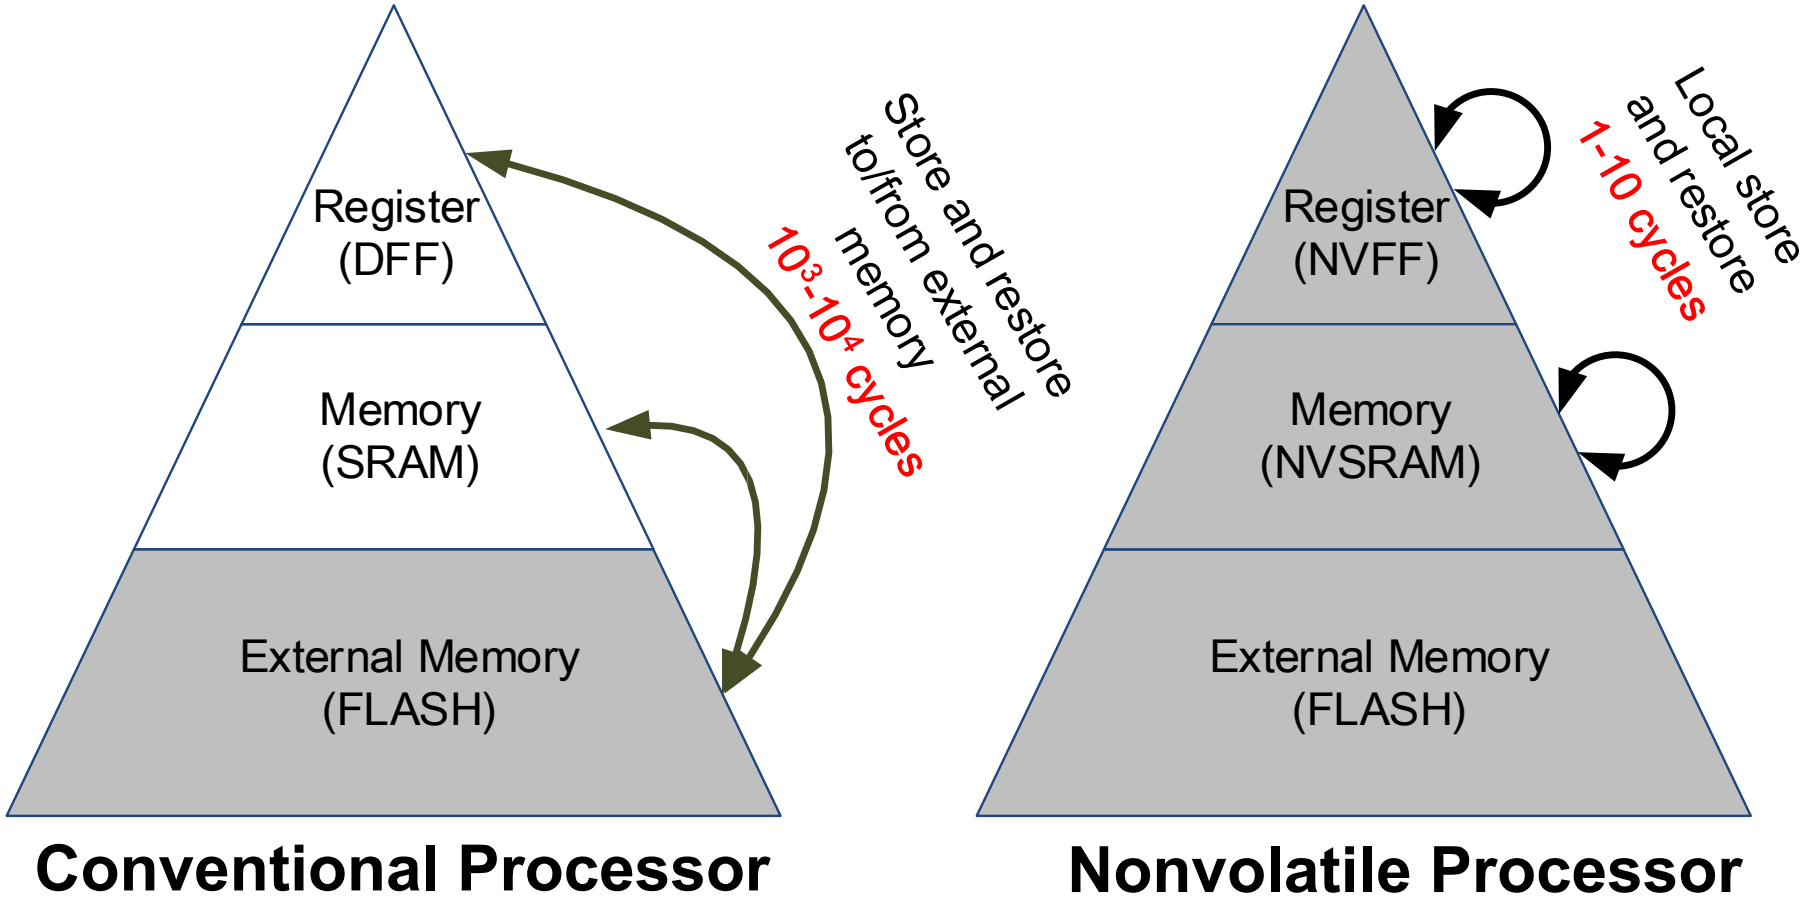

Source: DAC15 Invited

NVP Avoids Rollback with Faster Store/Restore!

Outline

- Background
- **Motivation**
- System Architecture
  - 1-Macro Structure
  - Time-Domain Adaption
  - Space-Domain Adaption
- Measurement & Analysis

FeFF-Based NVP

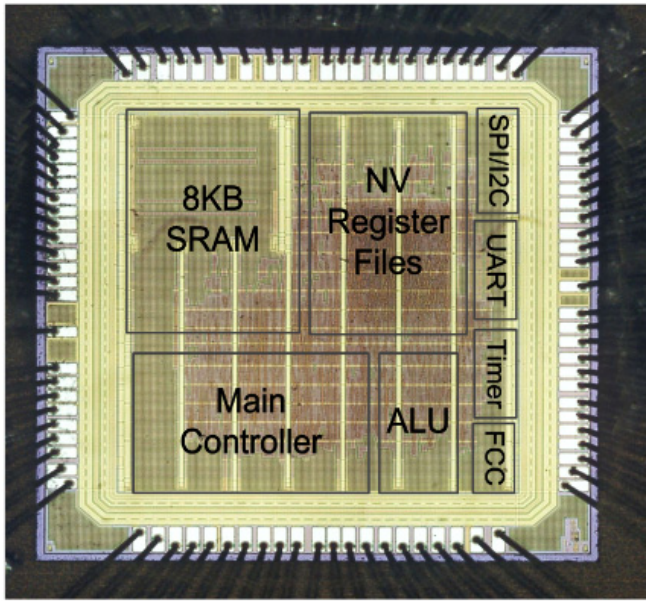

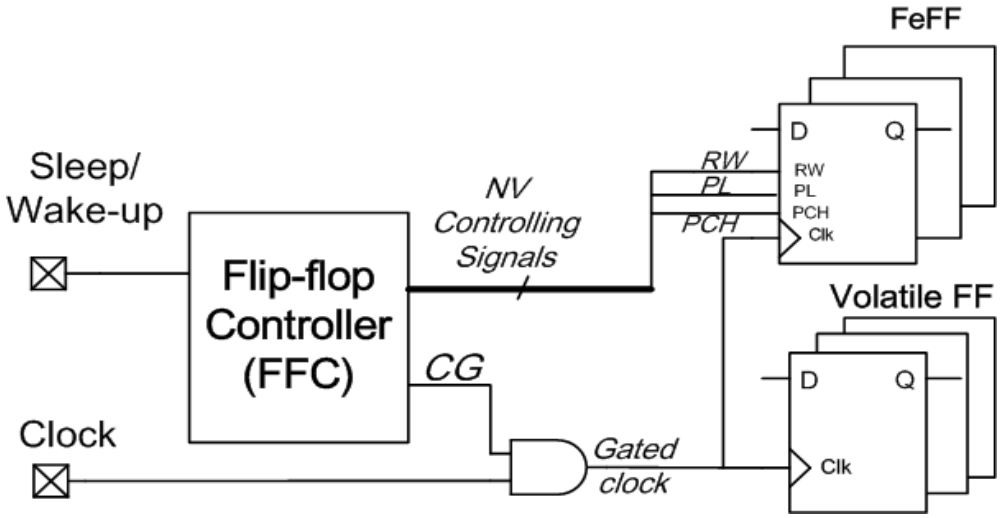

- **Full Parallel Arch.**
- **Pros: 2-4 OM Faster Store/Restore**
- **Cons: 90% Area Overhead**

Wang Y. et.al, ESSCIRC12

FeRAM Array Based NVP

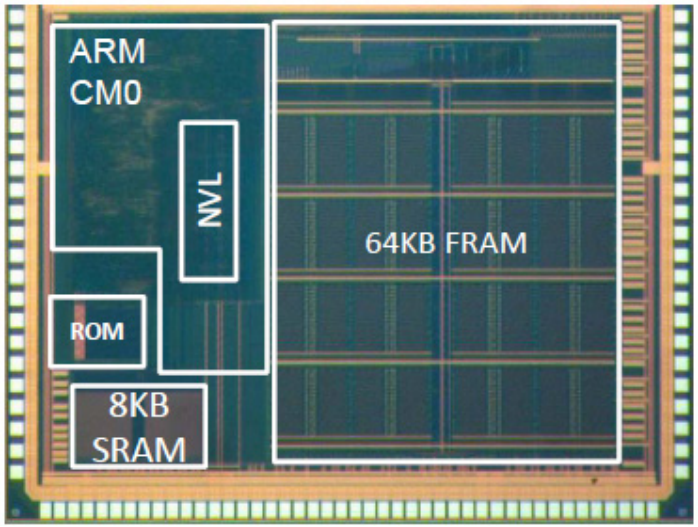

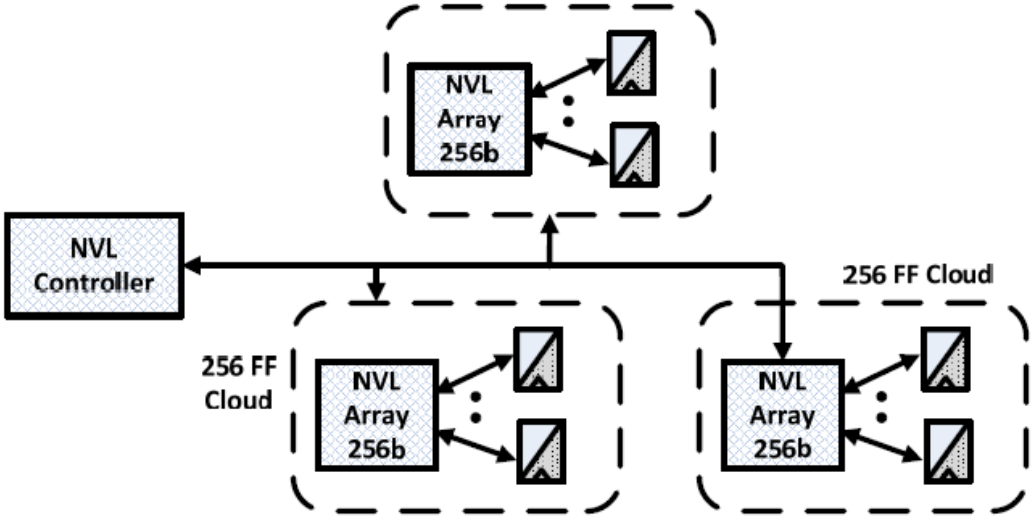

- **NVL Array Arch.**
- **Pros: Test & Area Friendly 3.6%**
- **Cons: Lower Speed 8MHz**

S. Bartling. et.al, ISSCC13

STT-RAM Based NVP

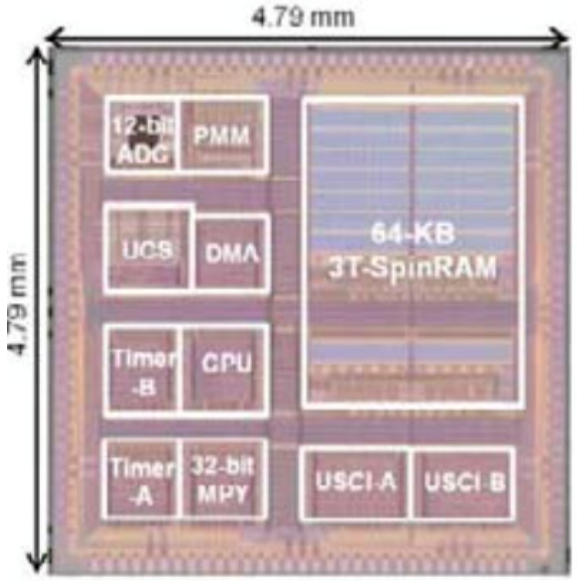

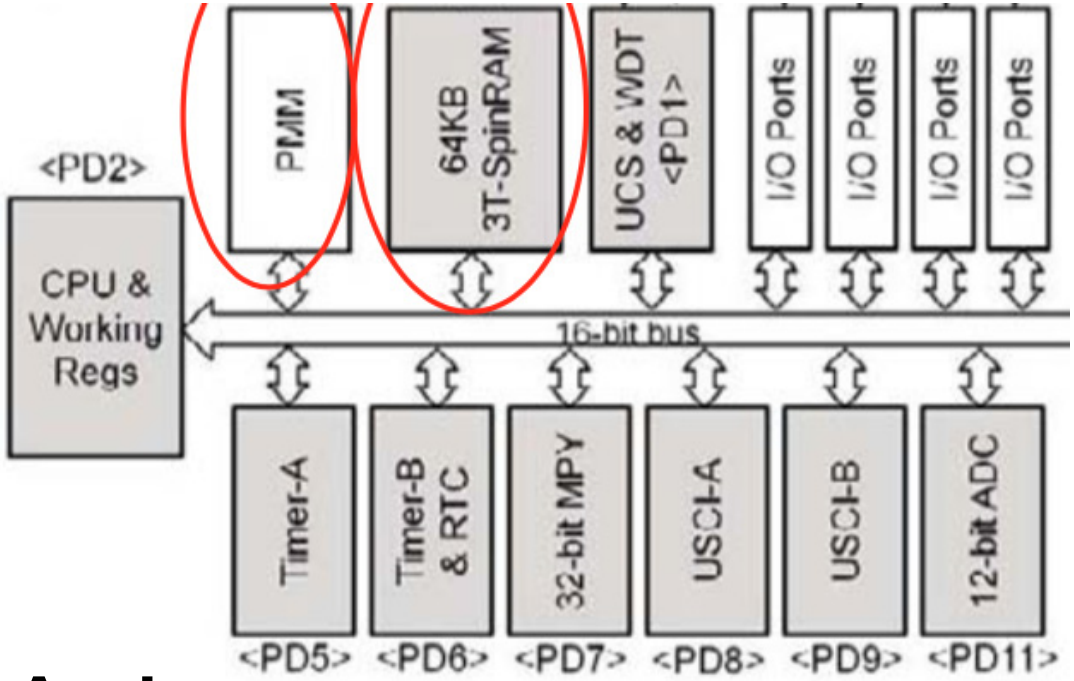

- **MFF Full Parallel Arch.**
- **Pros: 20MHz STT-RAM**
- **Cons: Frequent NVM Write** **N. Sakimura et al., ISSCC14**

Conventional NVP Architecture

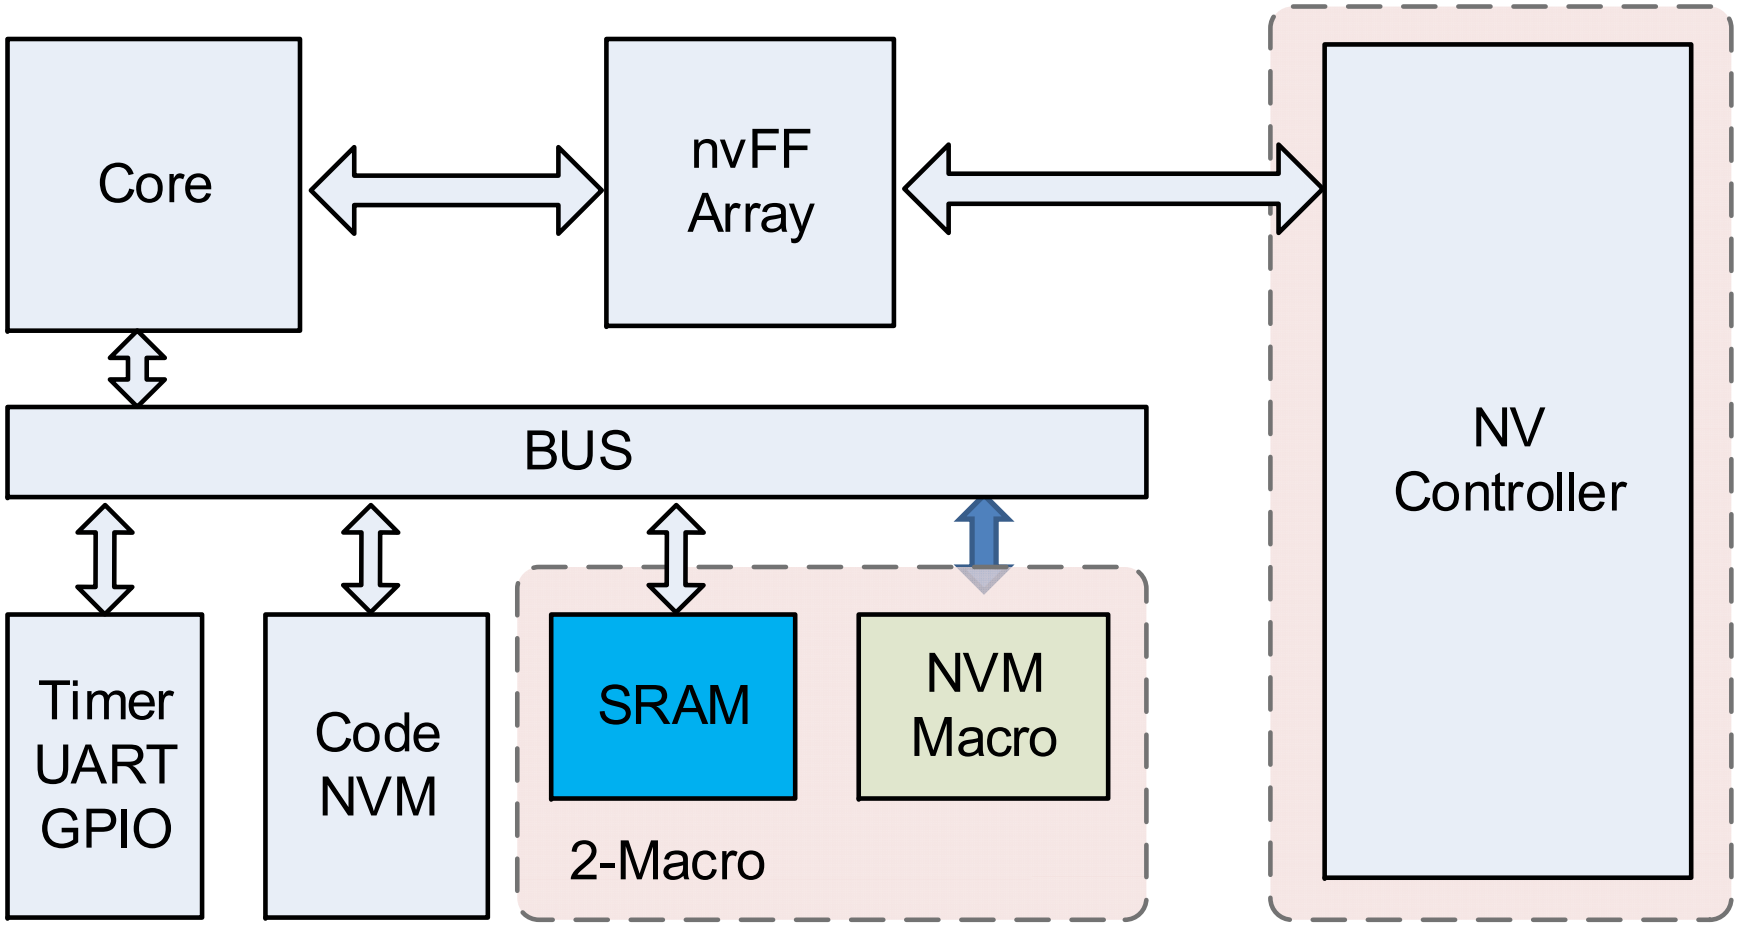

NVM Frequency & Bandwidth Limit

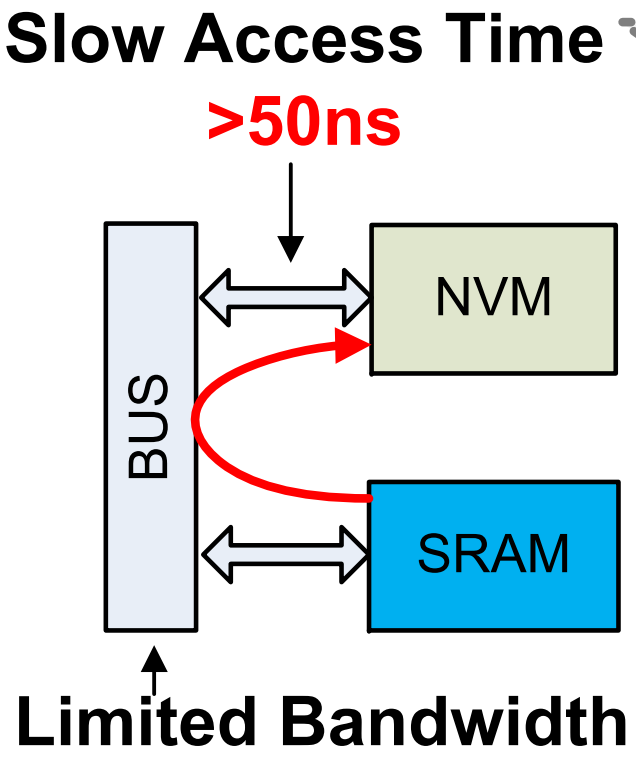

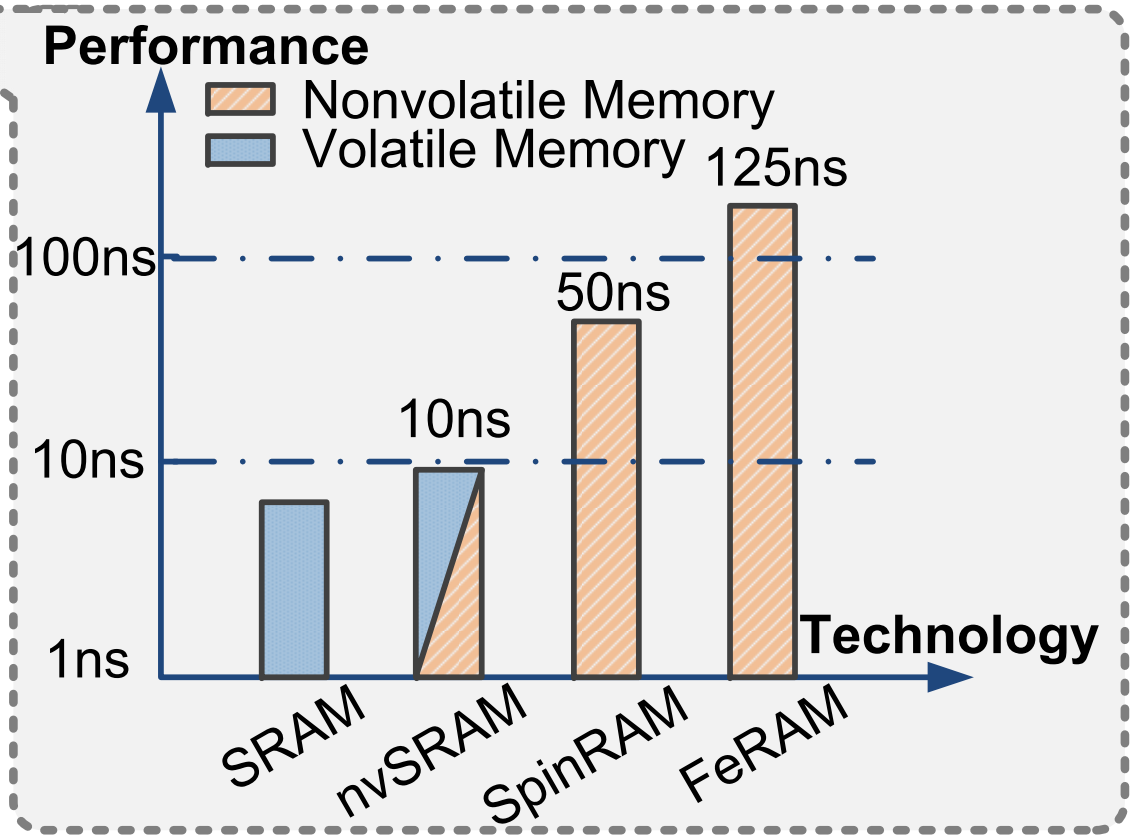

Limited bandwidth to store/restore, slow access to run.

Ignore Variable Power Off Req.

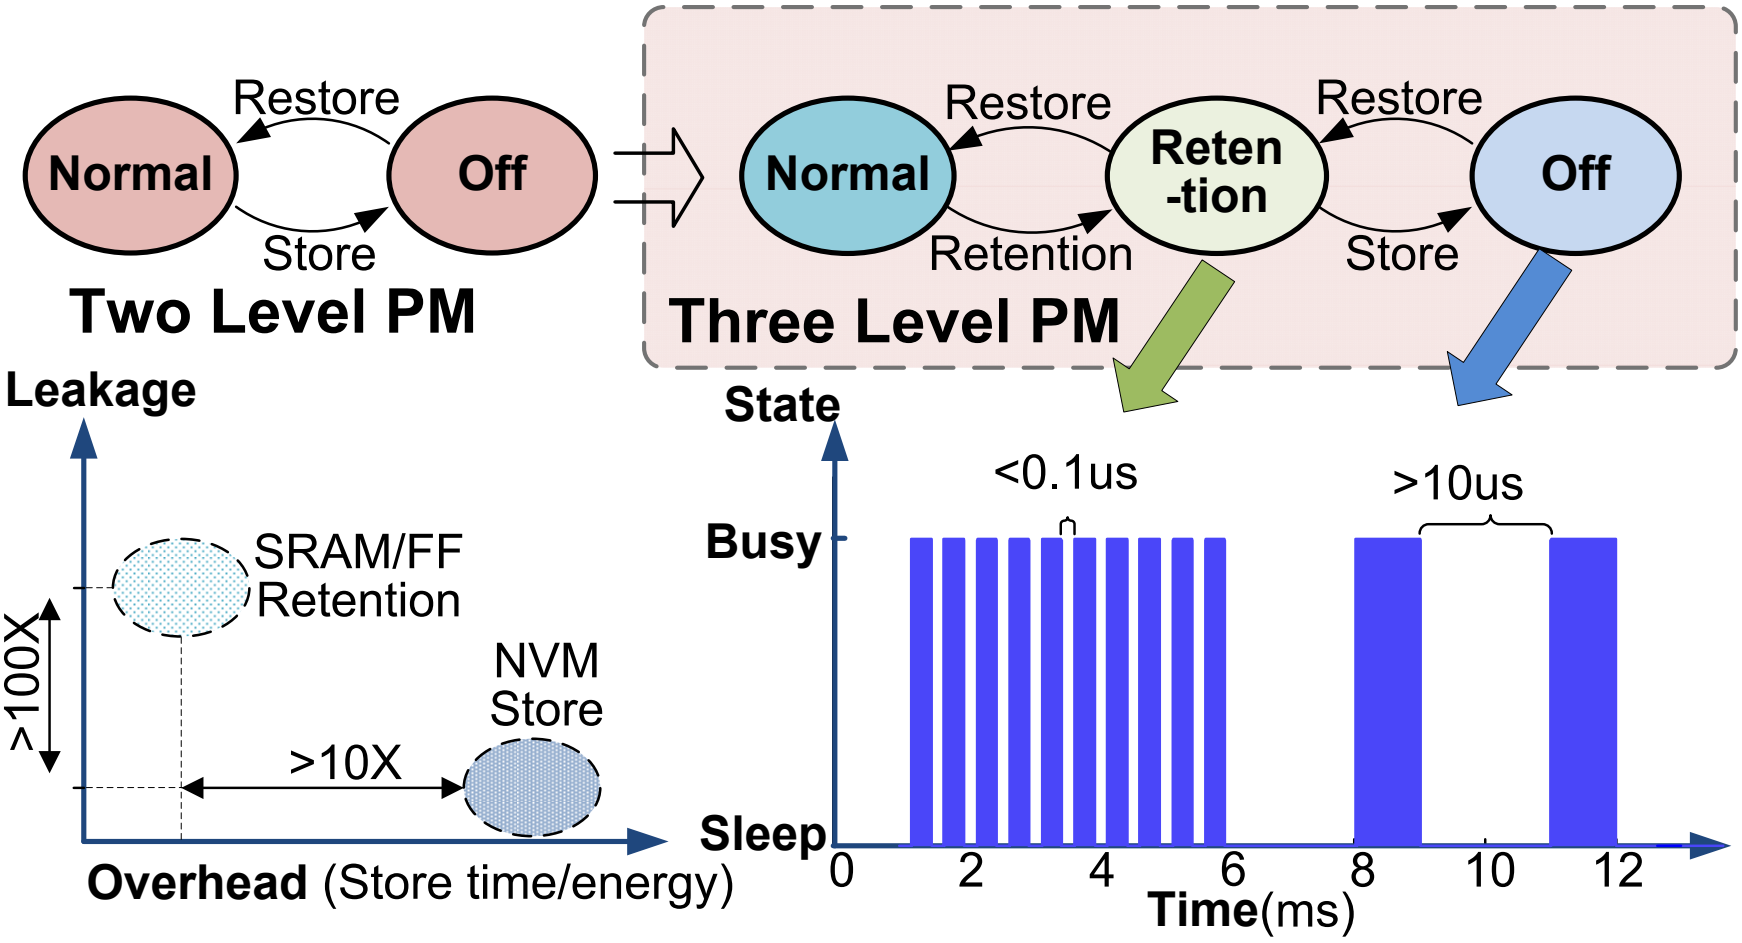

Breakeven point exists for FF retention vs NVM store.

Worst Case Design in NVM Backup

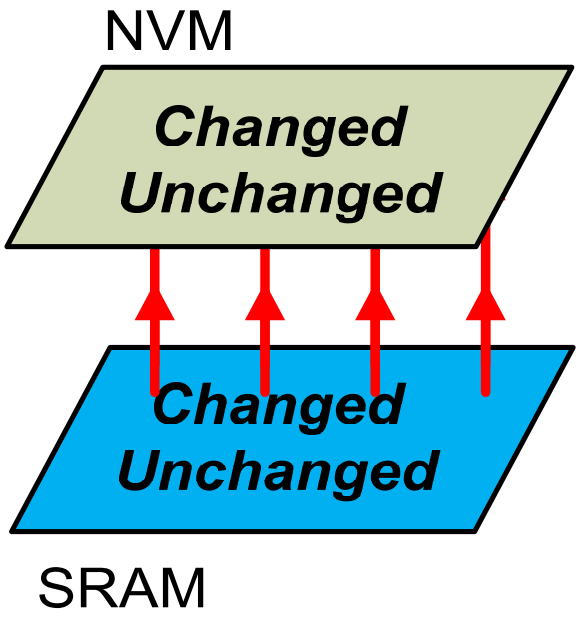

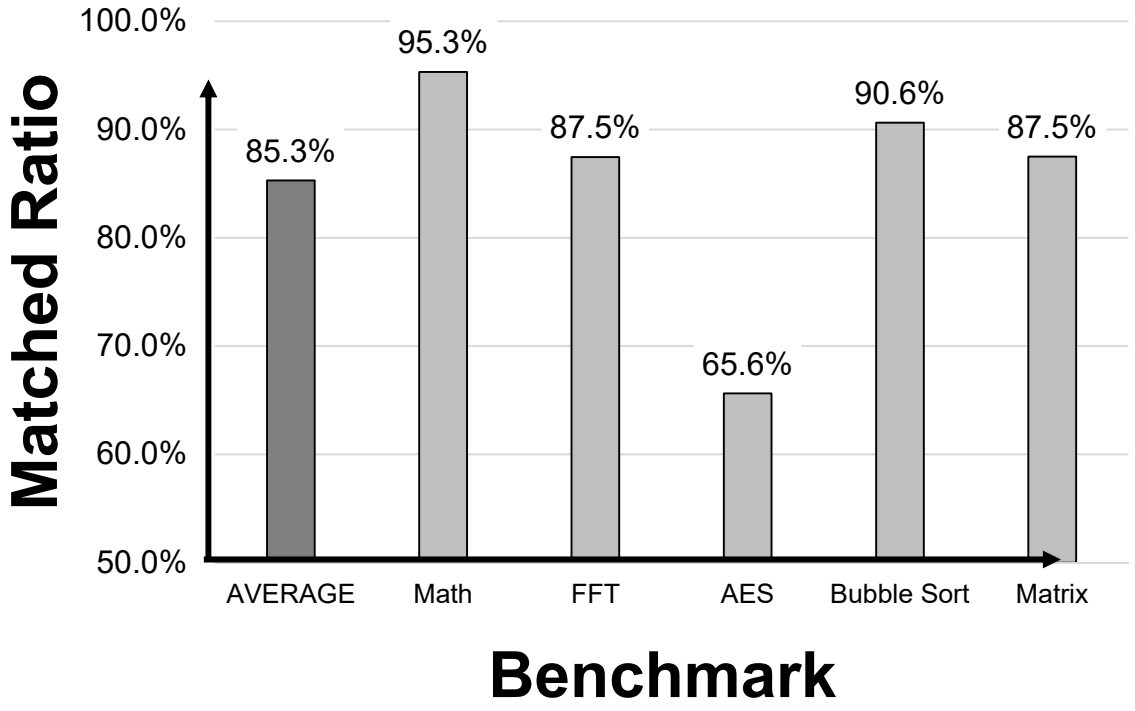

*Matched means data is unchanged.

Updating unchanged NVM wastes energy.

Outline

- Background
- Motivation
- **System Architecture**
  - 1-Macro Structure
  - Time-Domain Adaption
  - Space-Domain Adaption
- Measurement & Analysis

Proposed System Architecture

- 1-Macro adaptive memory
- Time-Domain adaption technique
- Space-Domain adaption technique

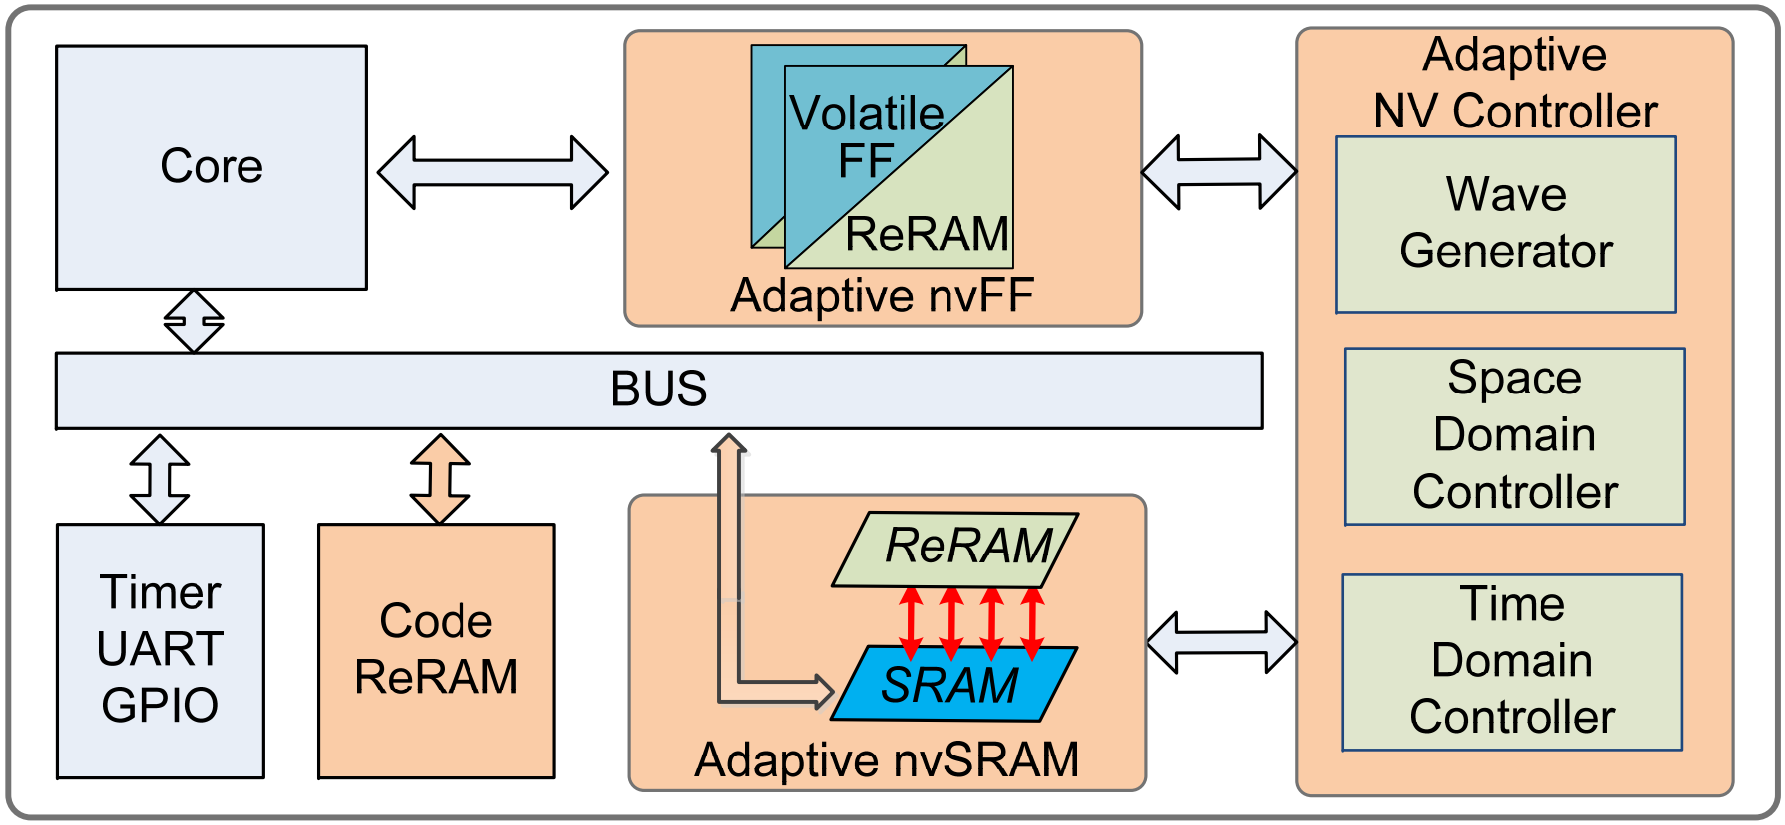

Fast & Parallel 1-Macro nvSRAM

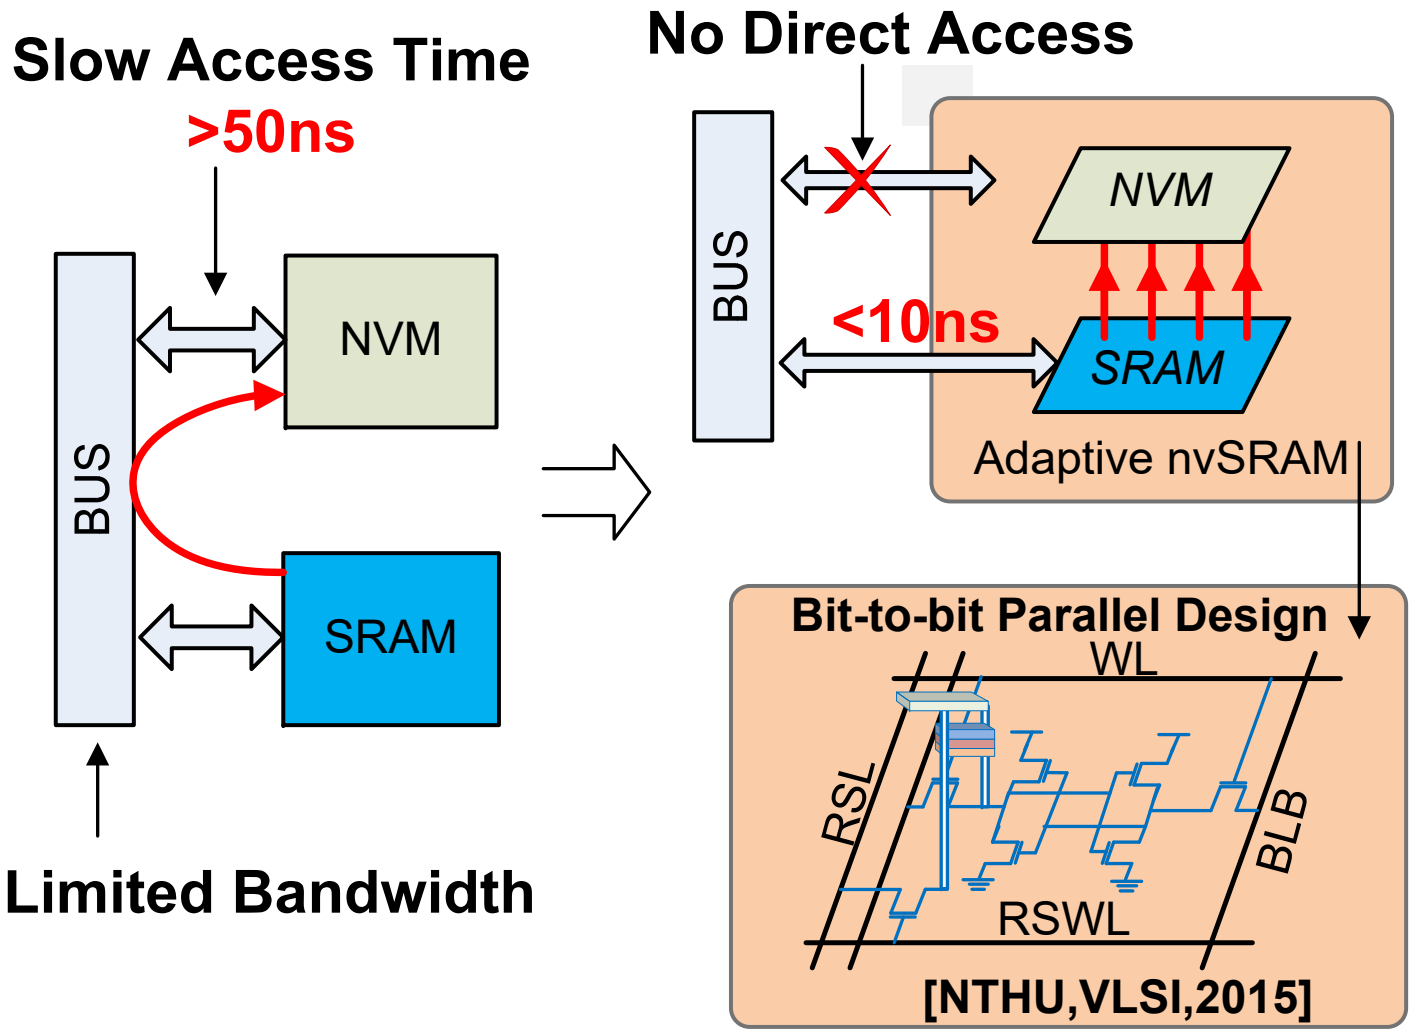

Serial (Many cycles) vs. Parallel (One cycle)

Time-Domain Adaptive Scheme

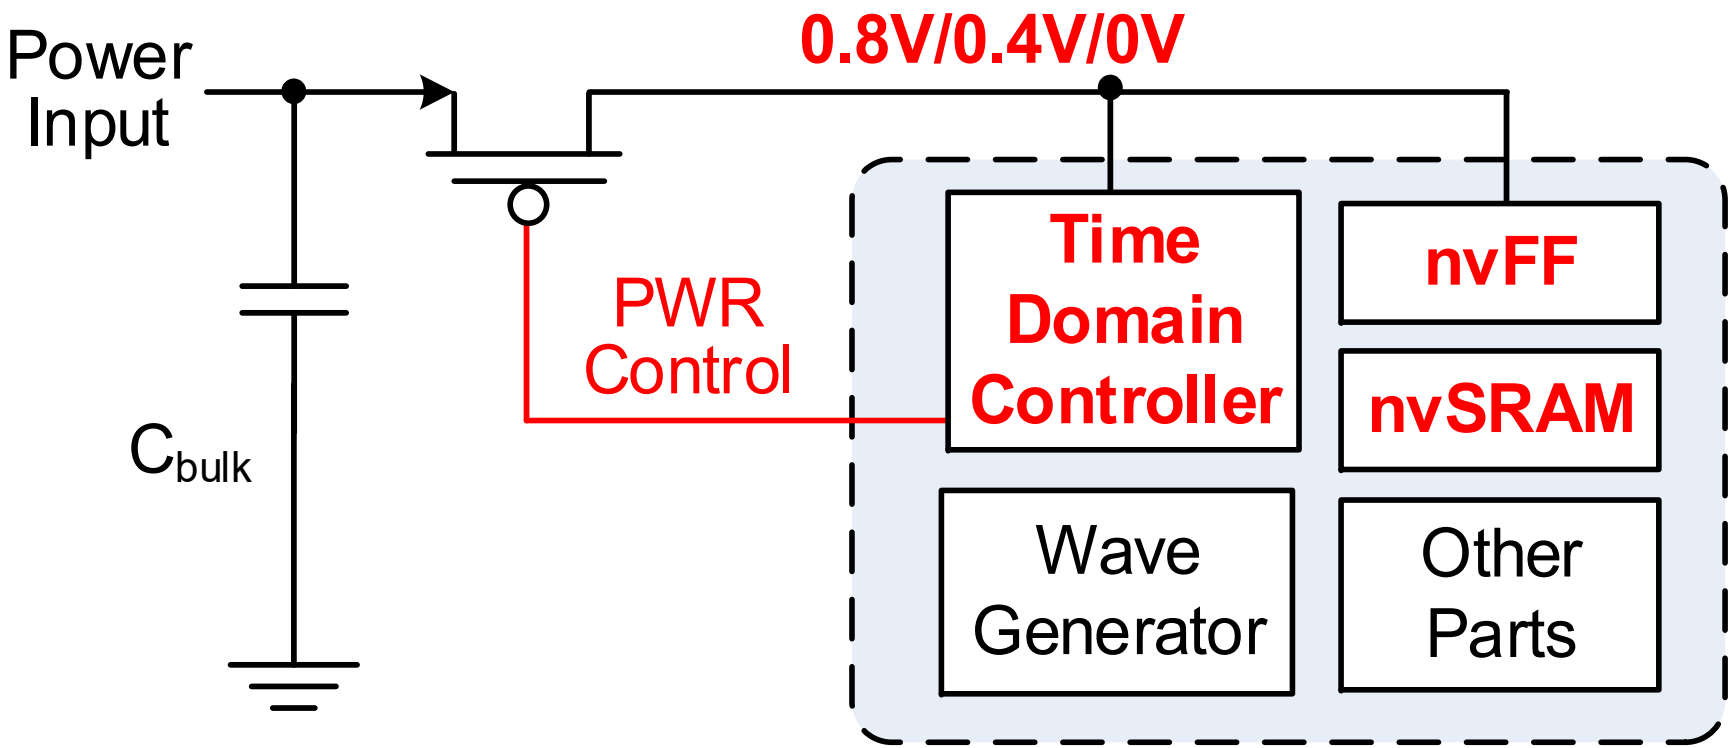

Power Modes: Normal/Retention/Off

Time-Domain NV Controller

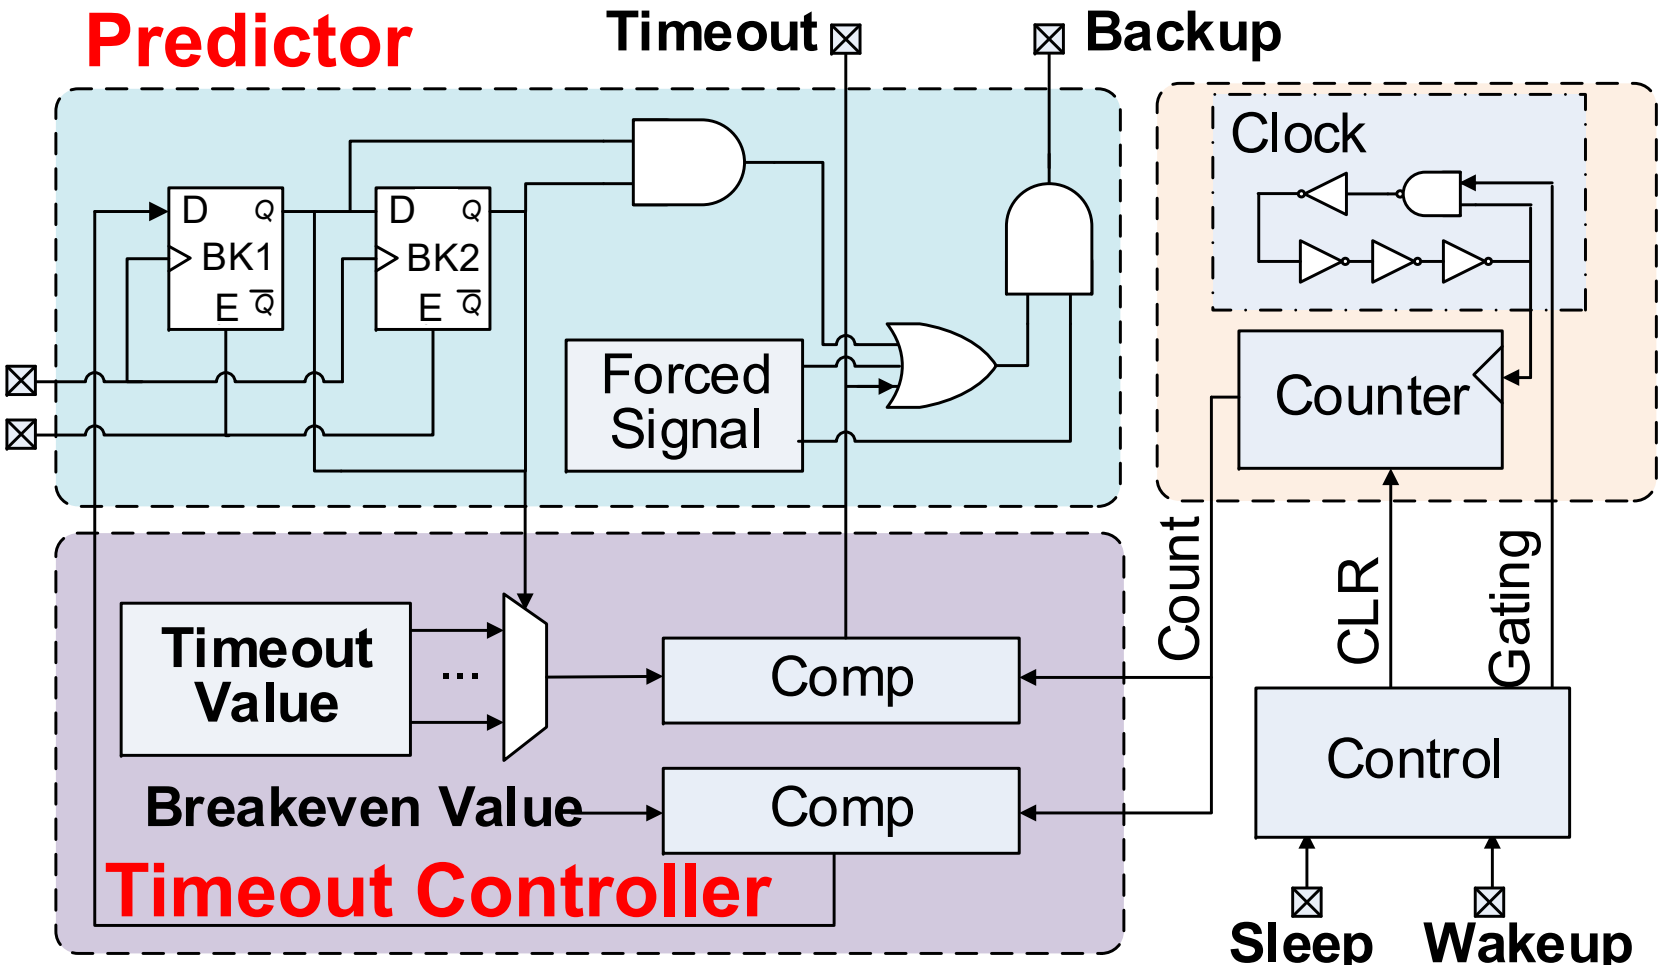

Retention or store, switch to store if timeout.

Time-Domain Adaptive nvFF

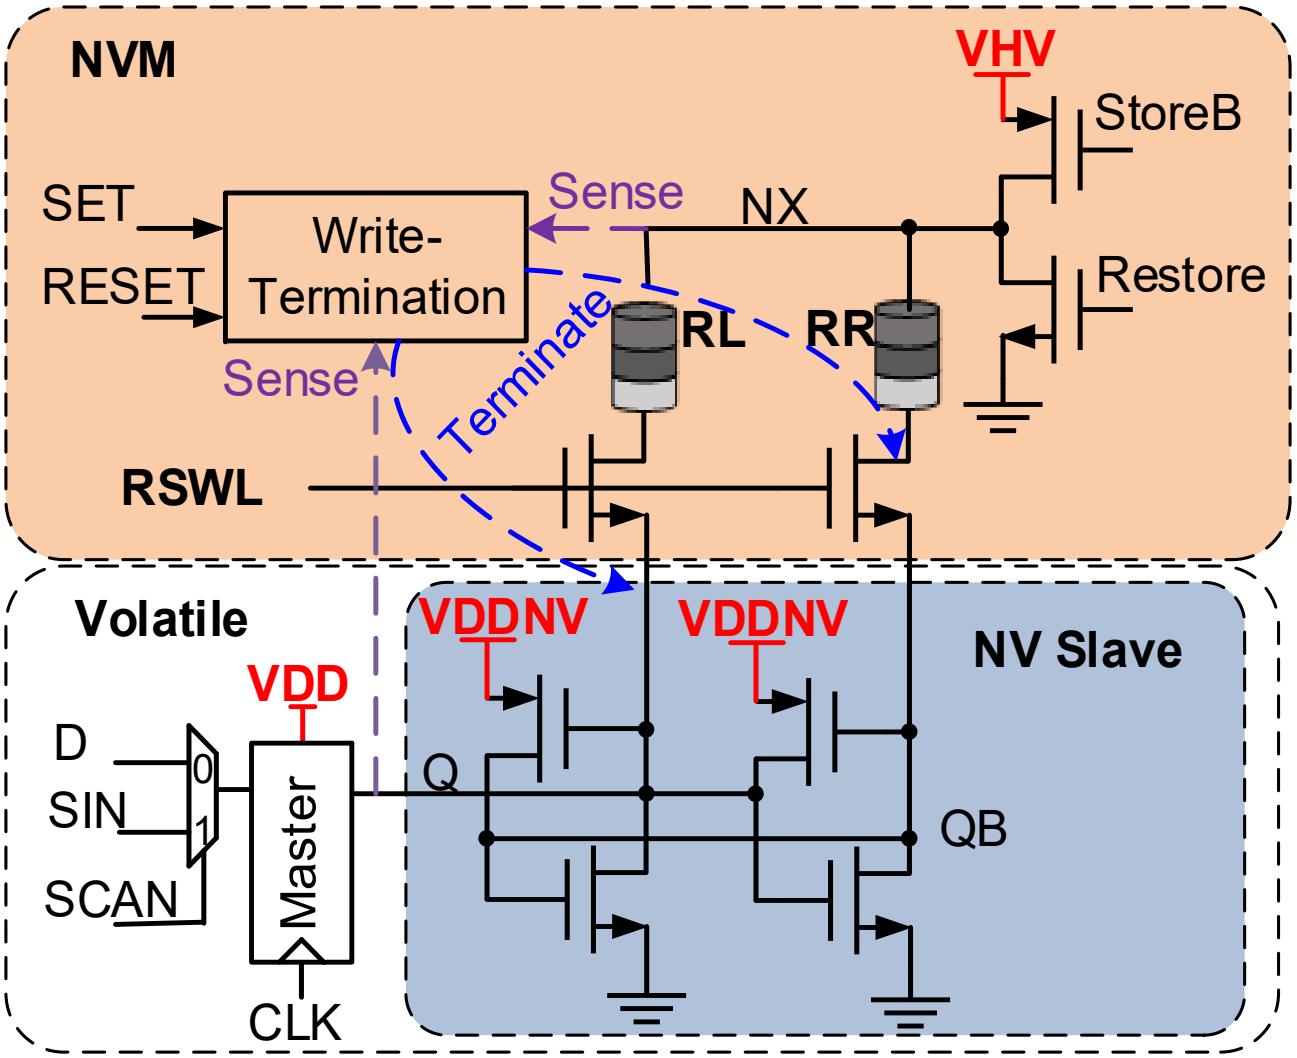

RSWL=0
VDD=0.8V
VDDNV=0.8V

Standard Flip-flop

Adaptive nvFF: Retention Mode

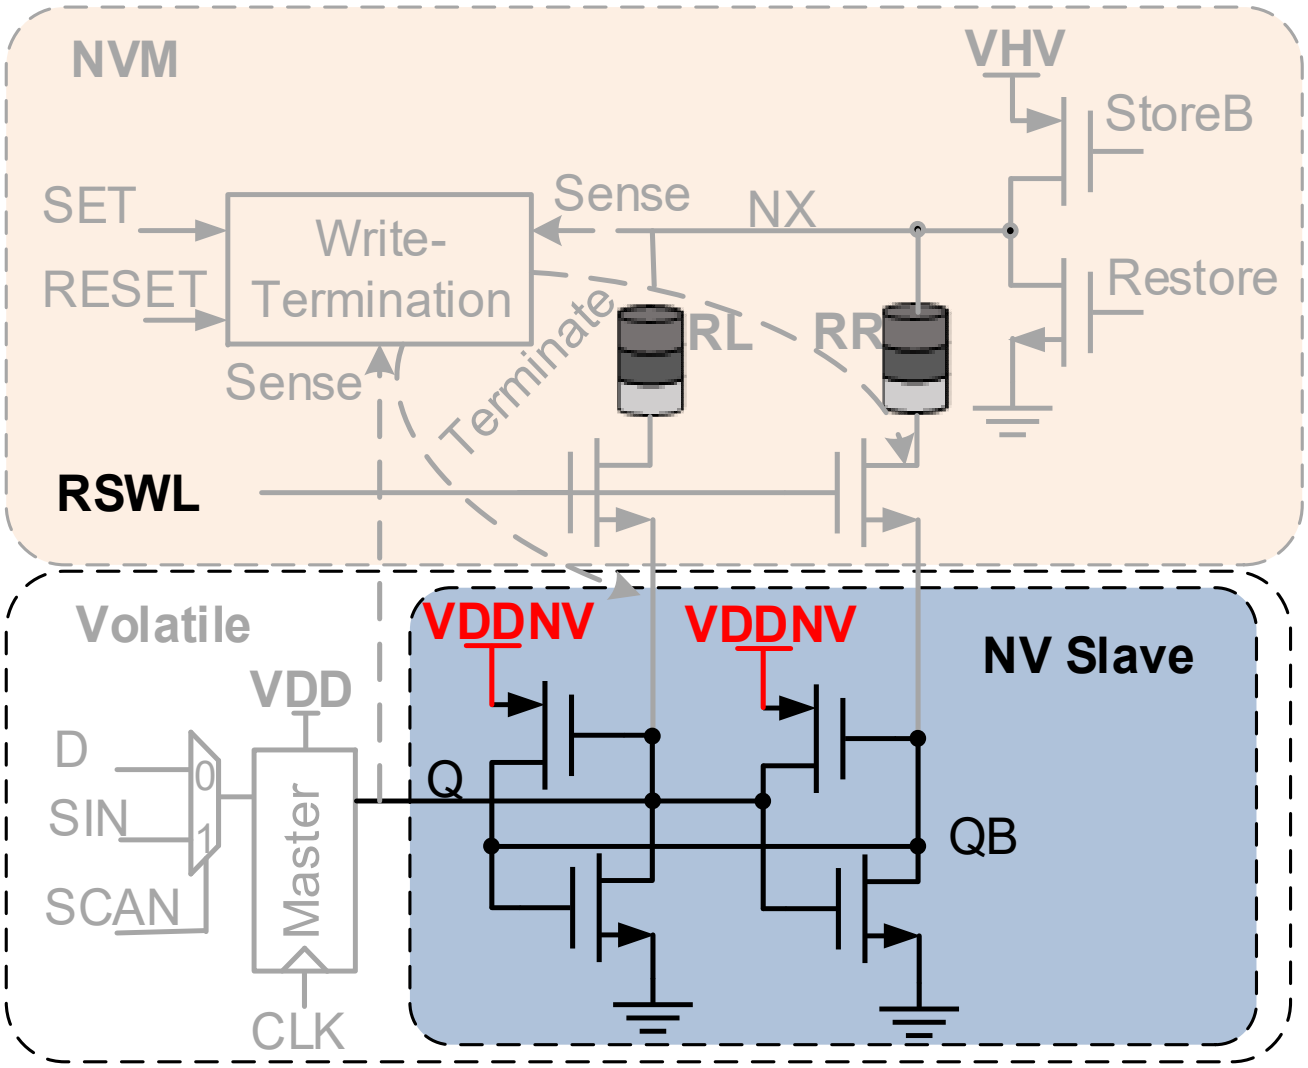

RSWL=0
VDD=0V
VDDNV=0.4V

Low Voltage Latch

Adaptive nvFF: Store/Restore

RSWL=1

VDD=0V

VDDNV=0.8V

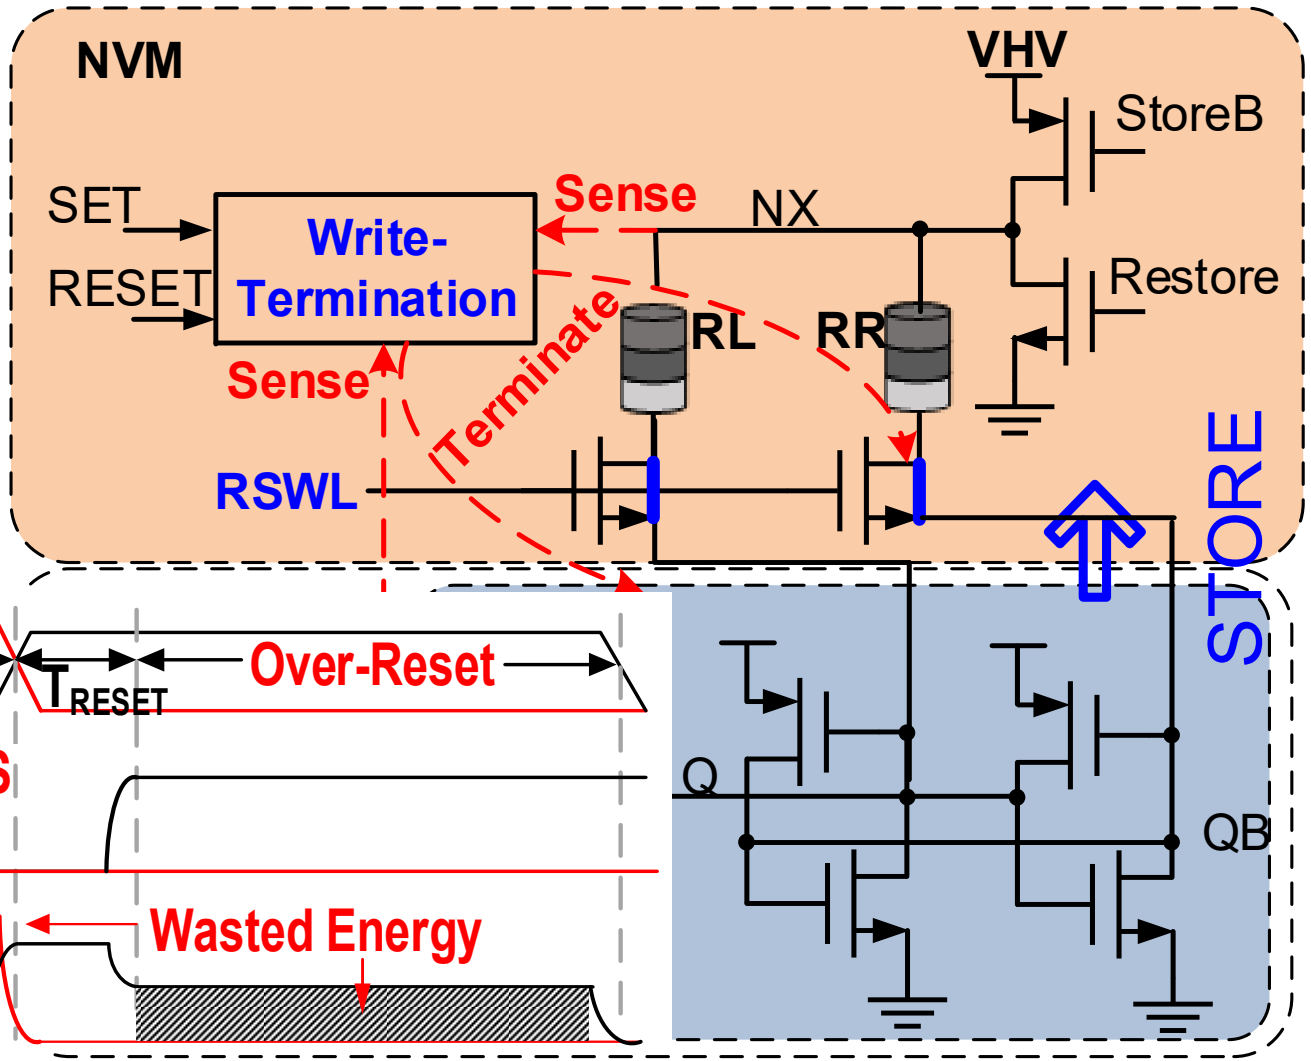

SWT reduces wasted energy for fast ReRAM cells.

Space-Domain Adaptive Scheme

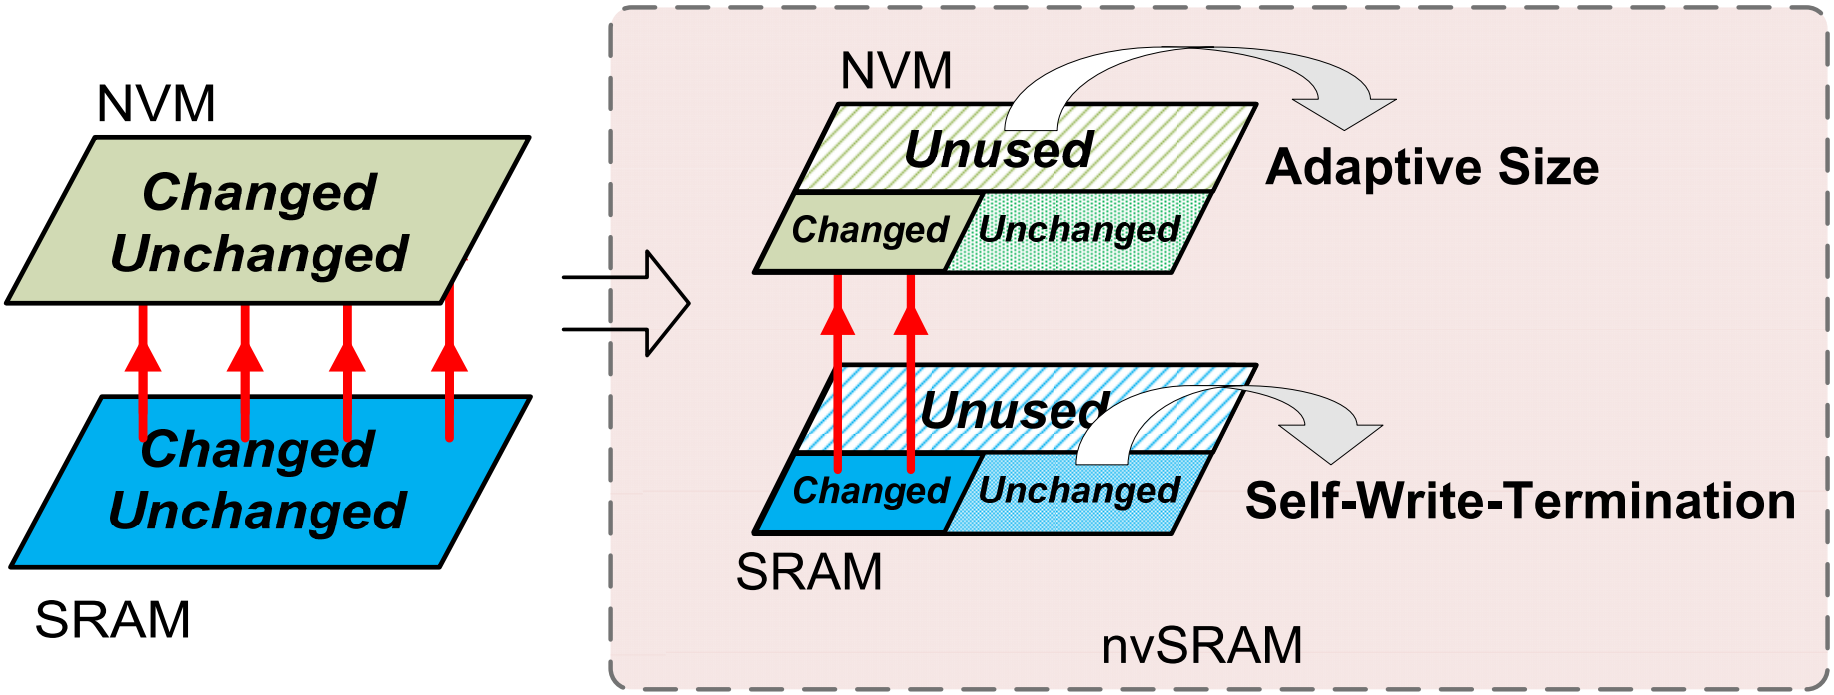

SWT avoids to overwrite unchanged data.
Adaptive nvSRAM size for low energy/ fast speed.

nvSRAM with Adaptive Scheme

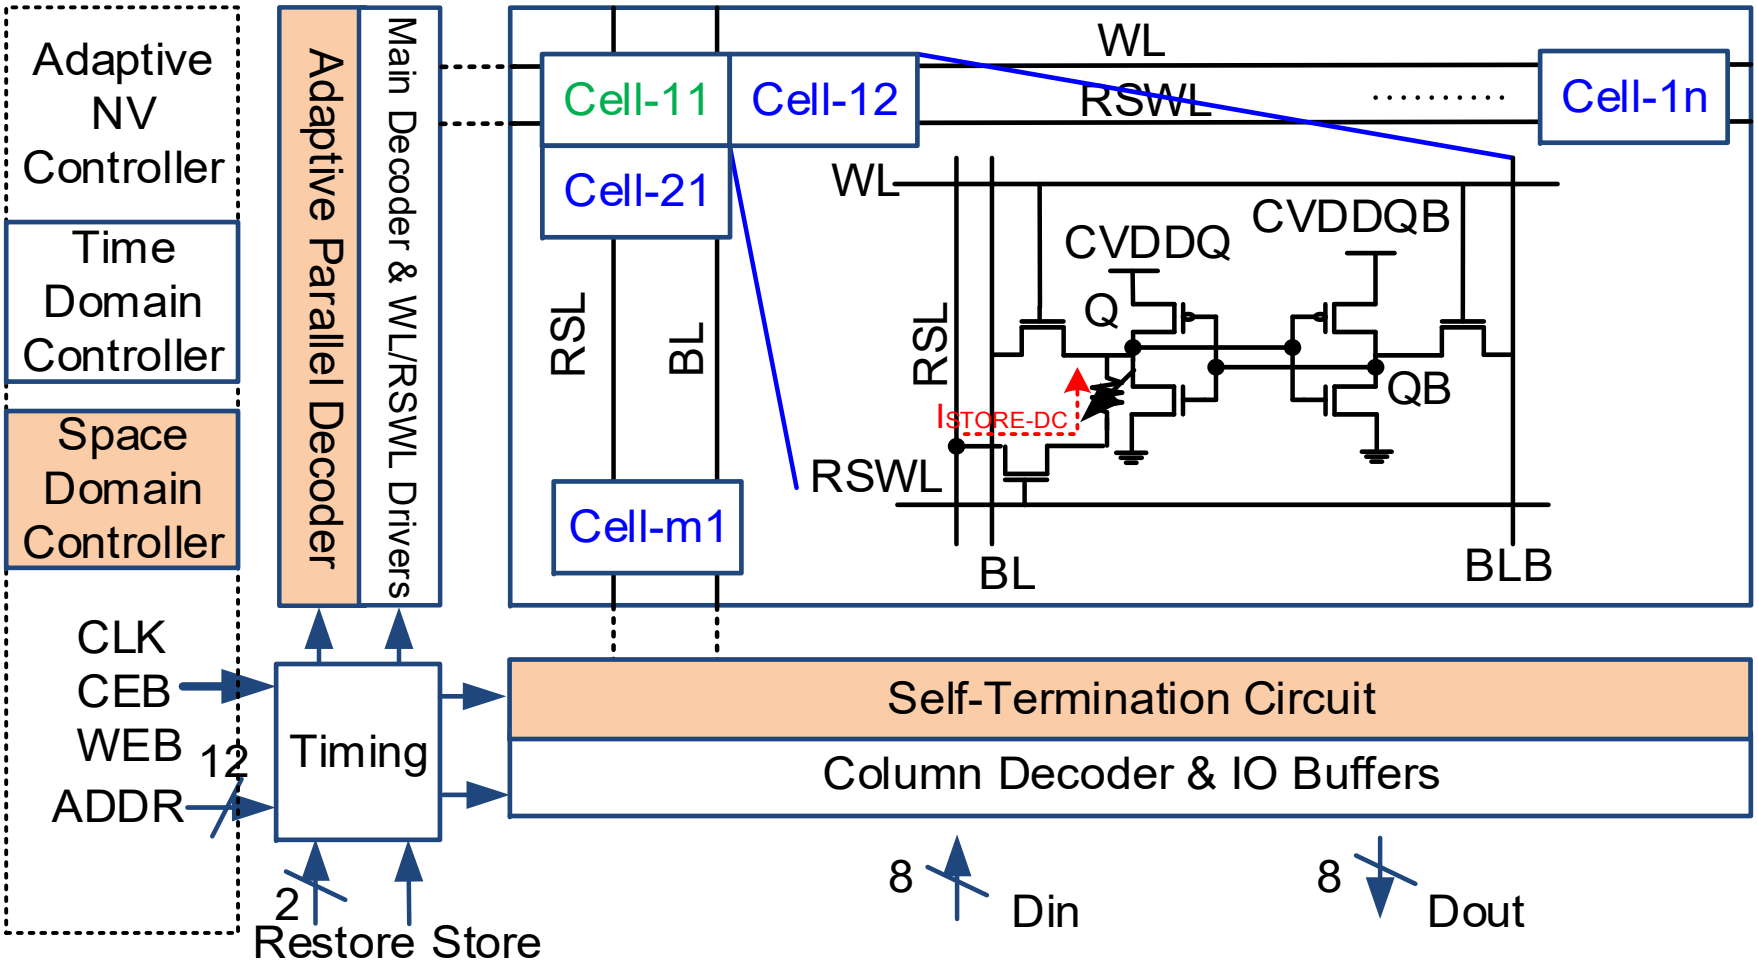

Space-Domain controller, Parallel decoder, SWT.

Space-Domain Controller

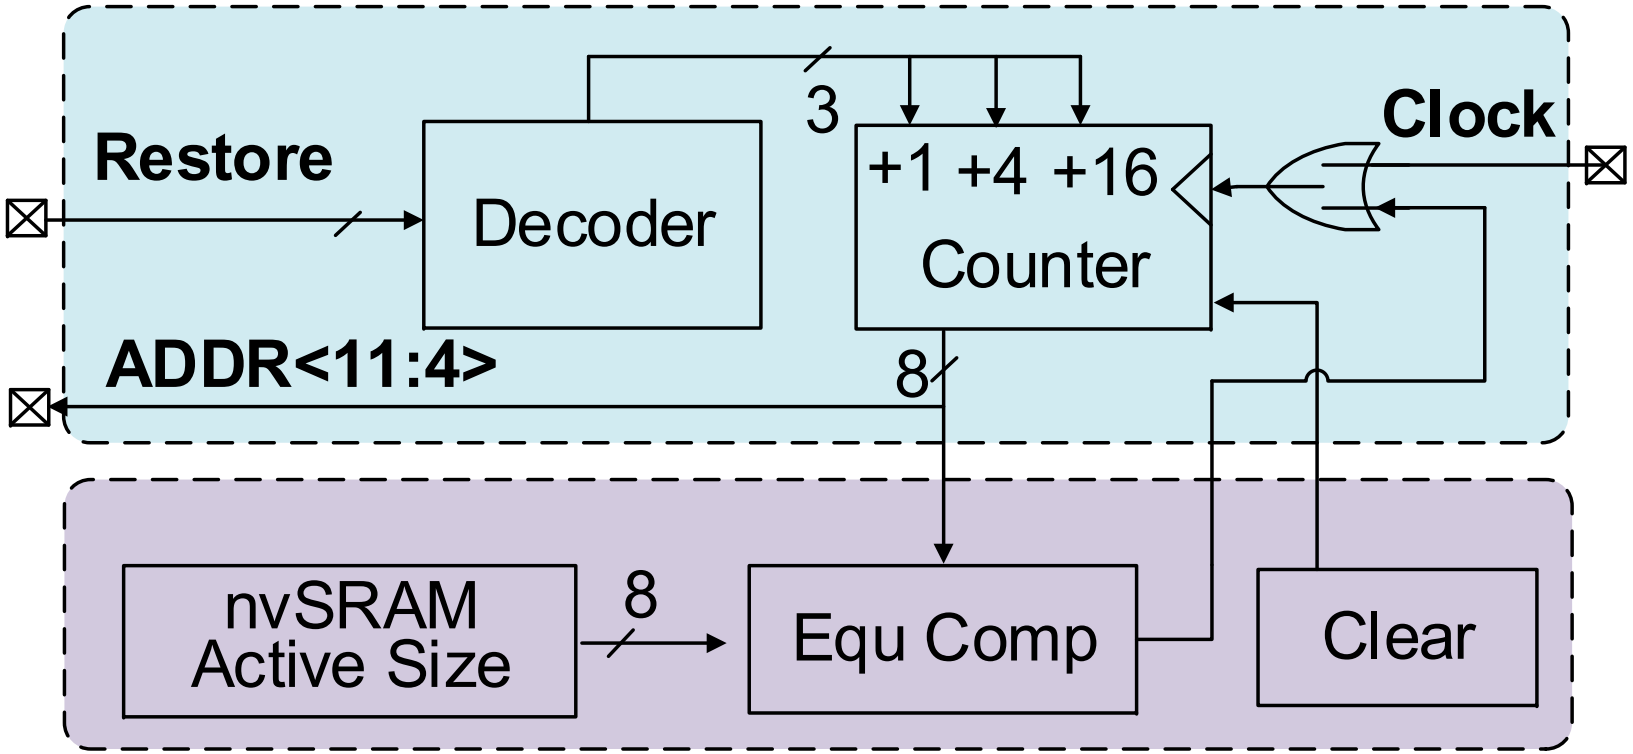

Address for different parallelism and memory size.

Adaptive Parallel Decoder

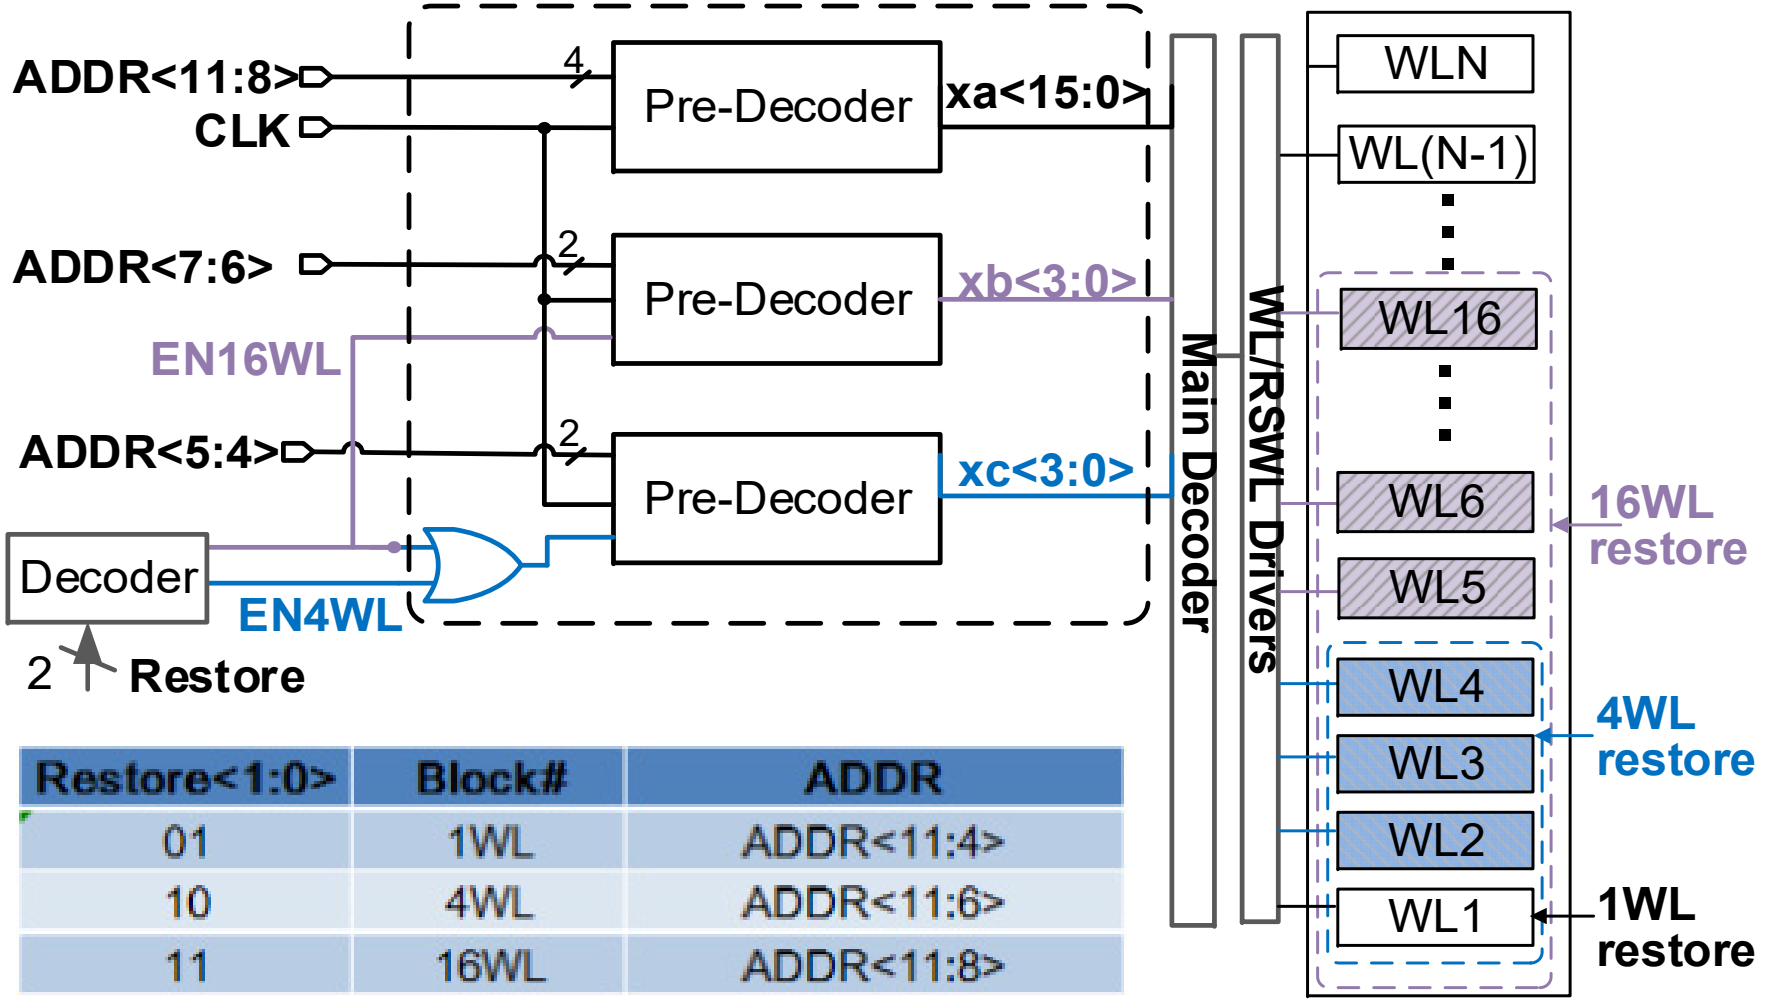

Restore<1:0>	Block#	ADDR
01	1WL	ADDR<11:4>
10	4WL	ADDR<11:6>
11	16WL	ADDR<11:8>

Different parallelism by combining xa/xb/xc.

Self-Write-Termination nvSRAM

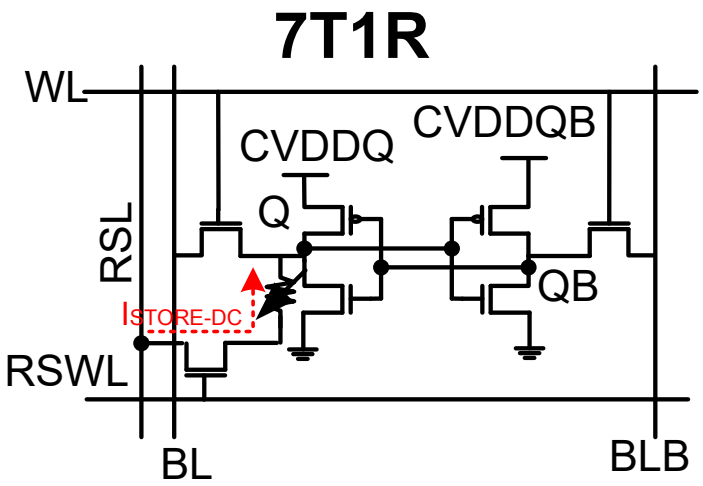

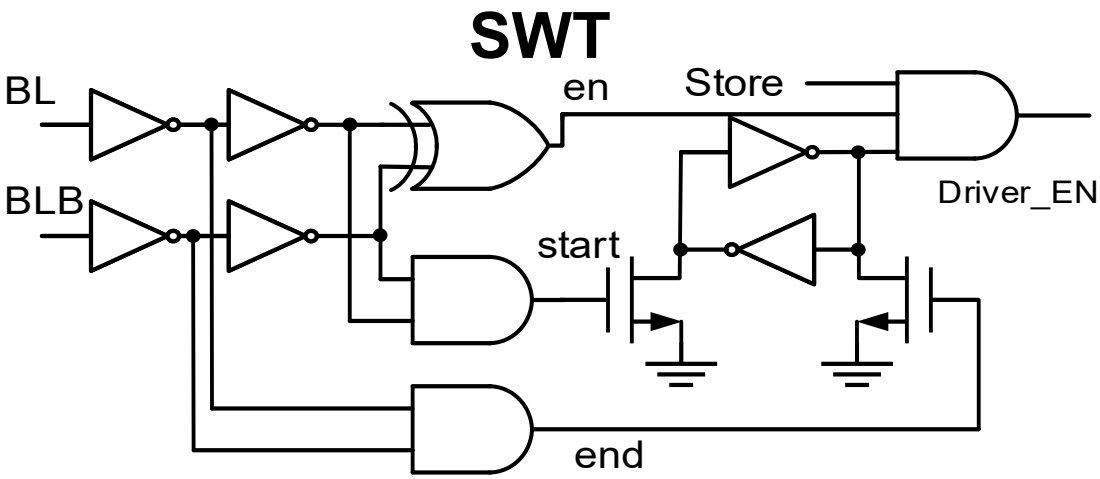

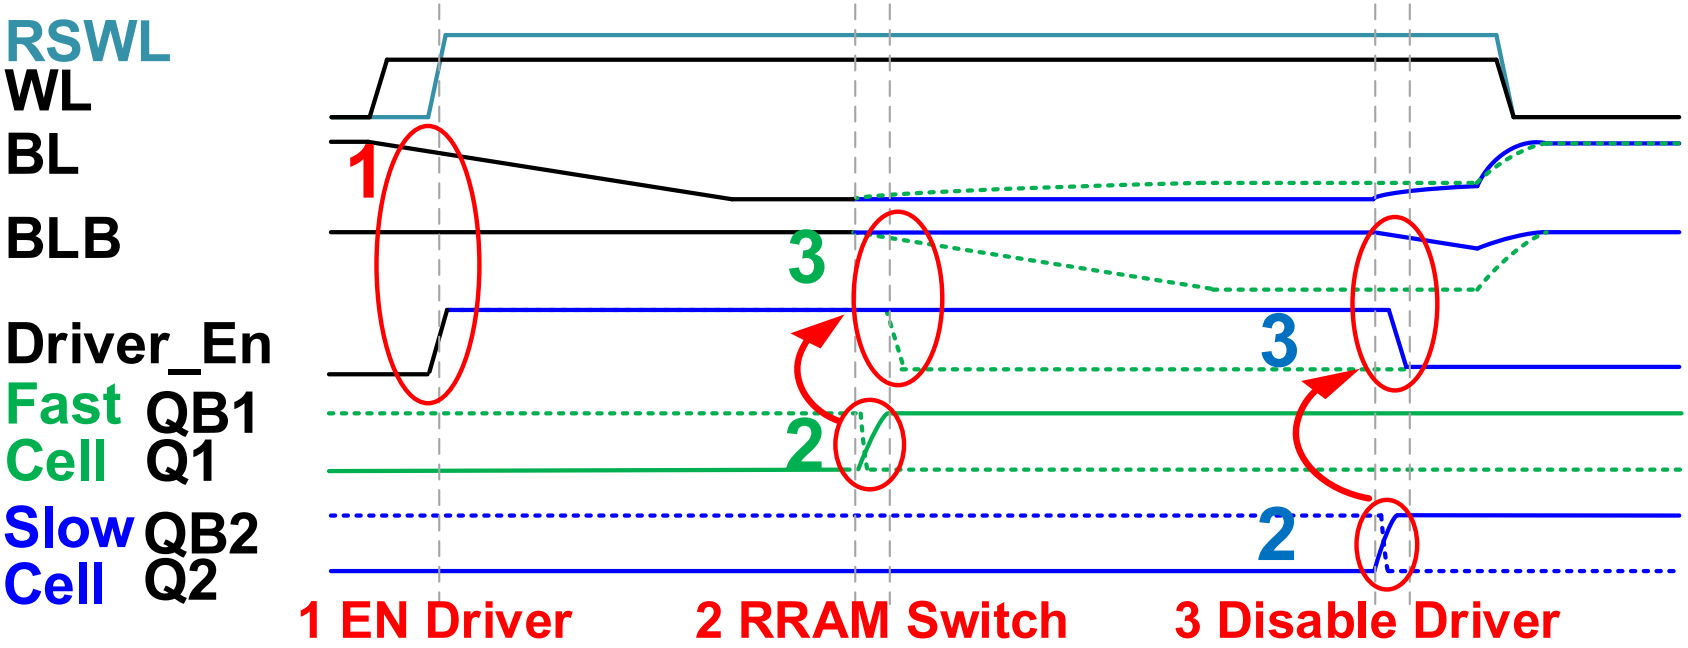

Outline

- Background
- Motivation
- System Architecture
  - 1-Macro Structure
  - Time-Domain Adaption
  - Space-Domain Adaption
- **Measurement & Analysis**

ReRAM Based NVP & Metric

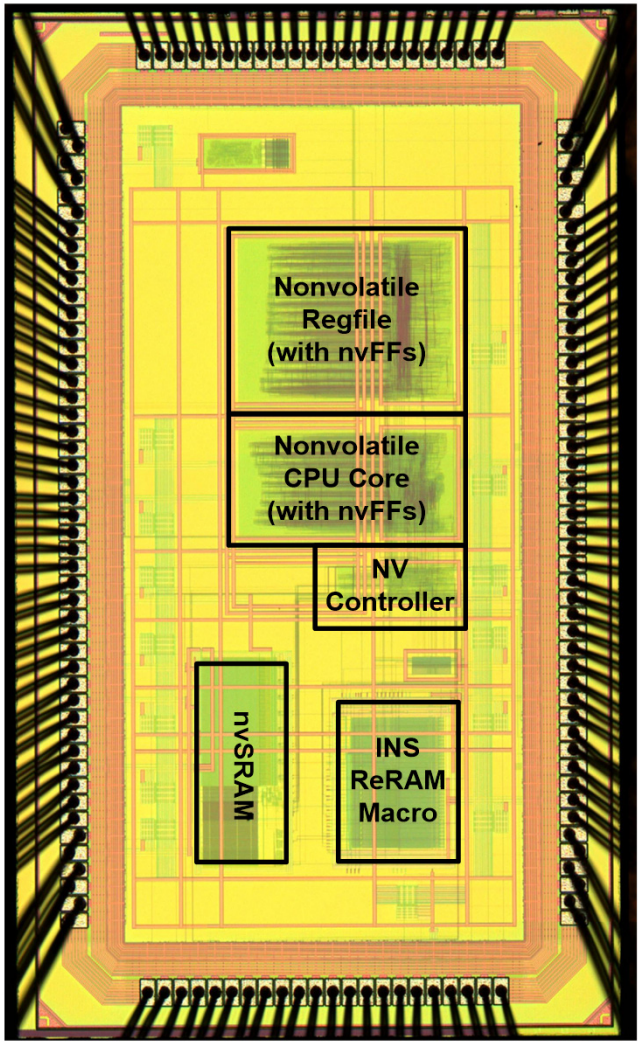

Technology	65nm SVT CMOS +ReRAM
NVM	8KB ReRAM 4KB nvSRAM 1422bits nvFF
Supply	0.8V(Core), 3V(HV)
Chip Area	1560 x 2860 um^2
Nonvolatile Area	nvFFs: 5.98% nvSRAM: 4.71% ReRAM Macro: 5.84%
Frequency	>100 MHz
System Restore Time	20 - 170 ns
System Restore Energy	0.45nJ (Avg)
System Store Time	4us- 1.02 ms
System Store Energy	0.40uJ (Avg)

1-Macro NVM Arch@100MHz

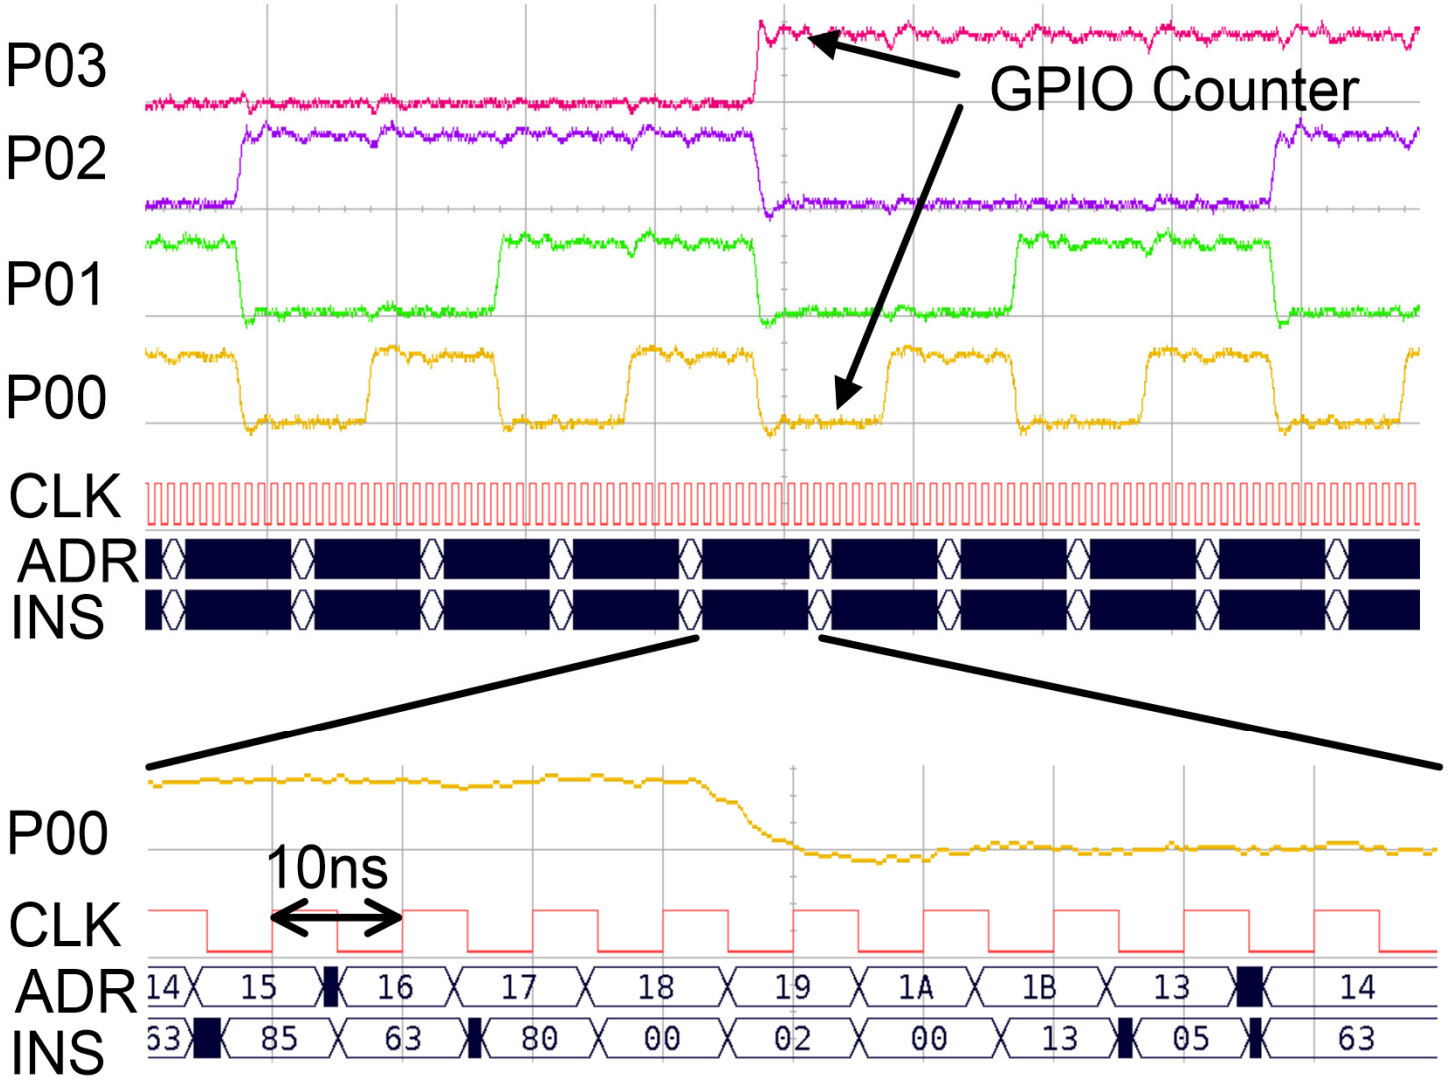

Time-Domain Adaptive: Retention

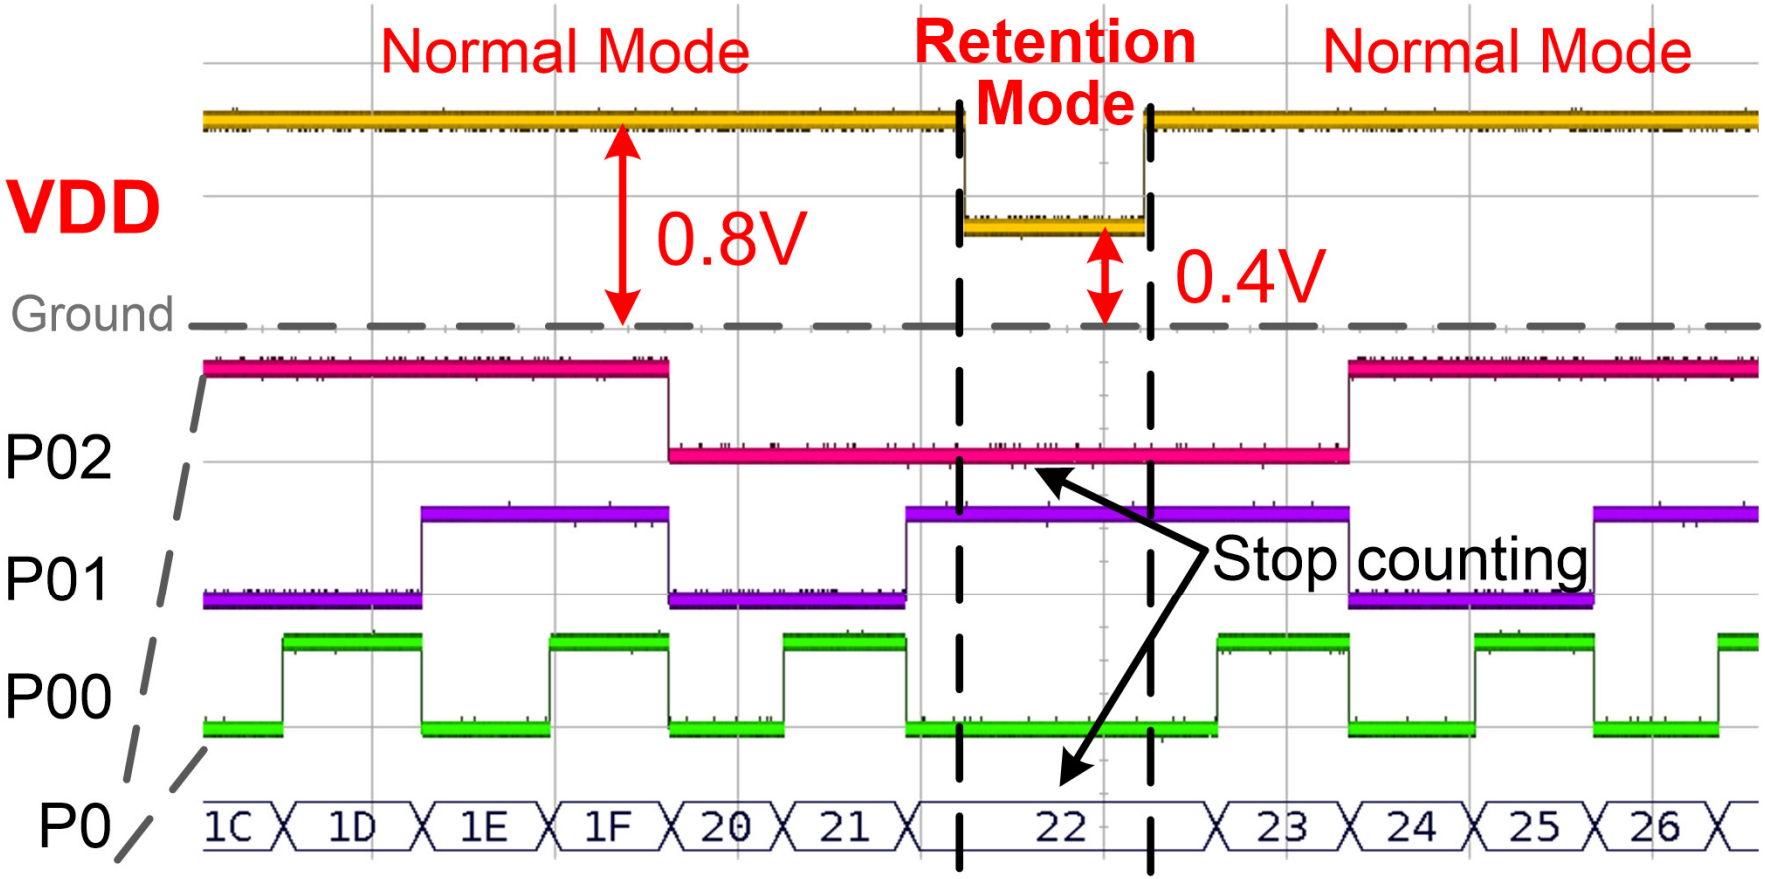

Time-Domain Adaptive: Restore

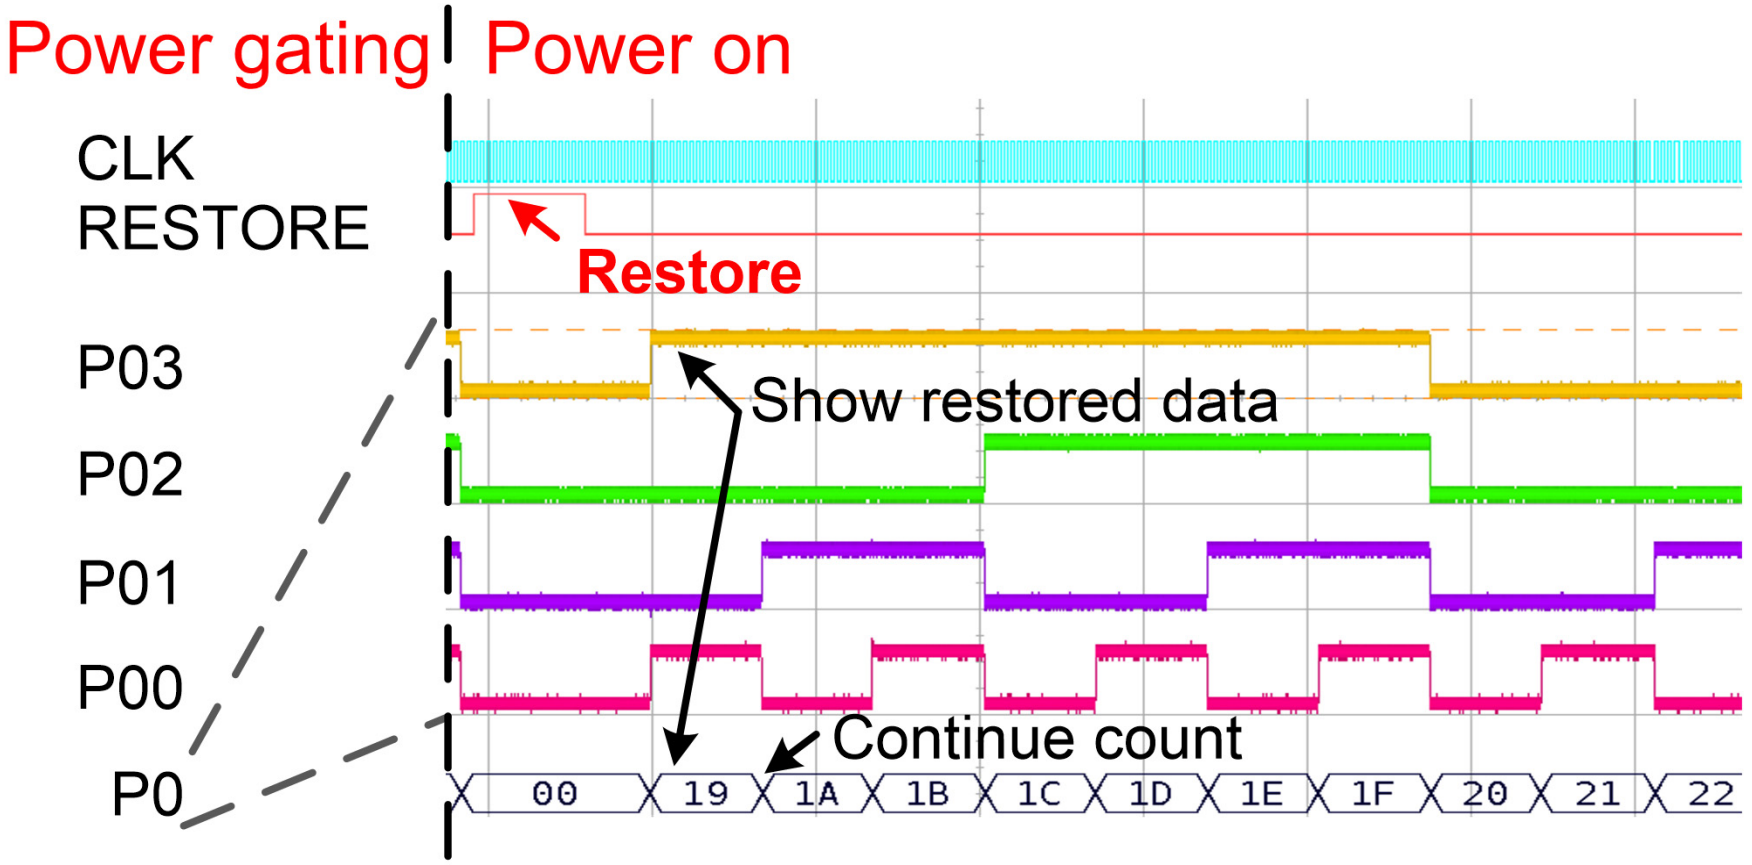

Space-Domain Adaptive nvSRAM

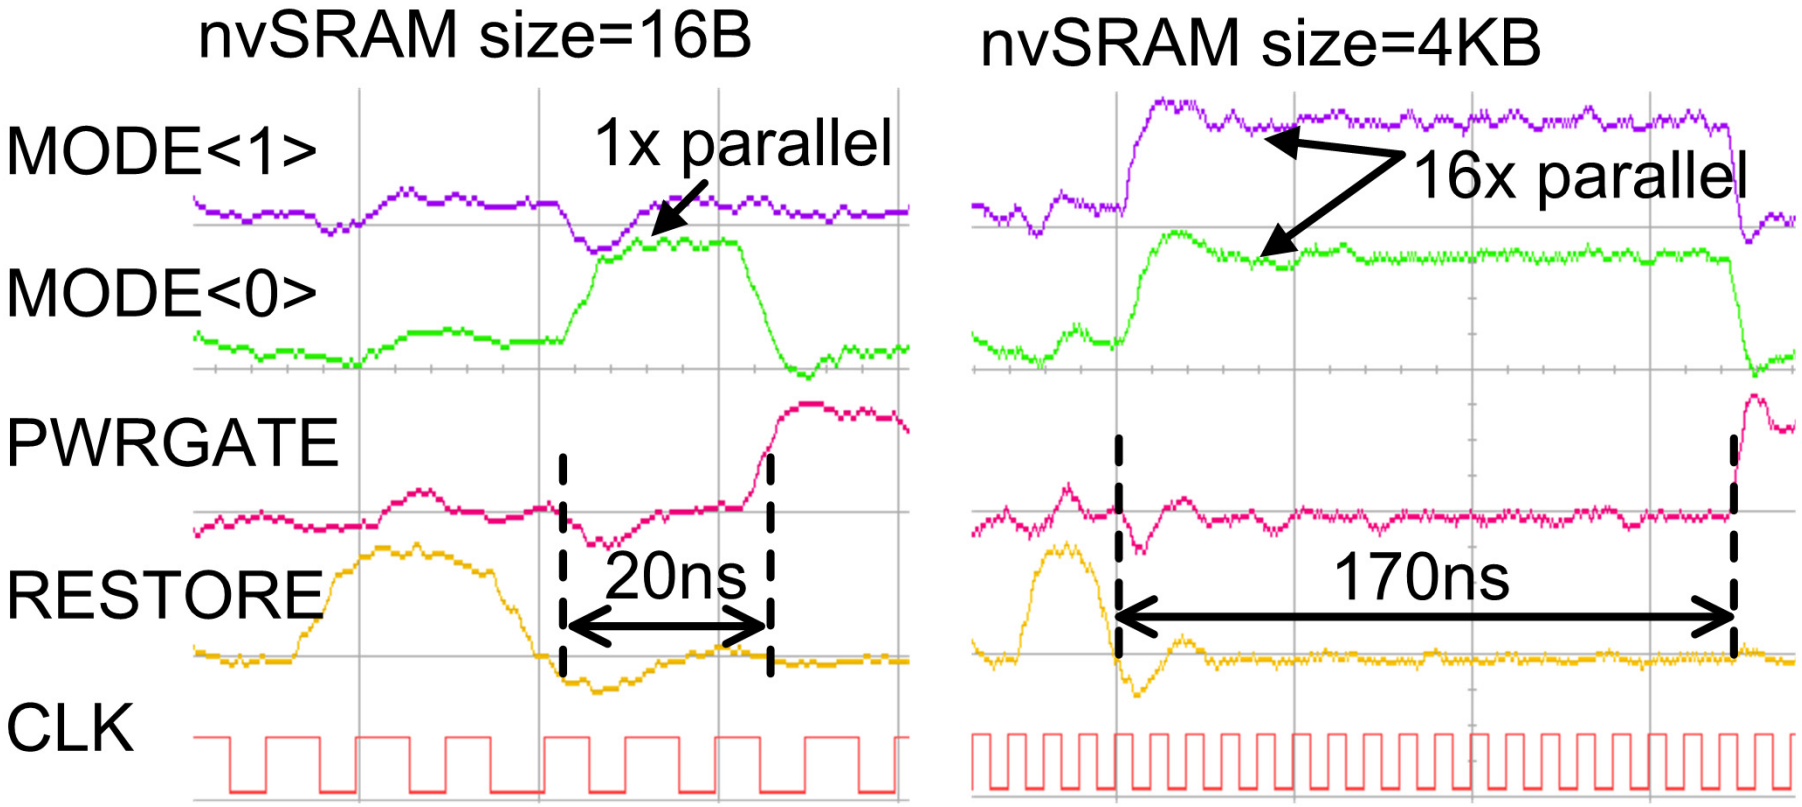

Energy/Time Savings Breakdown

Store Energy Reduction Methods

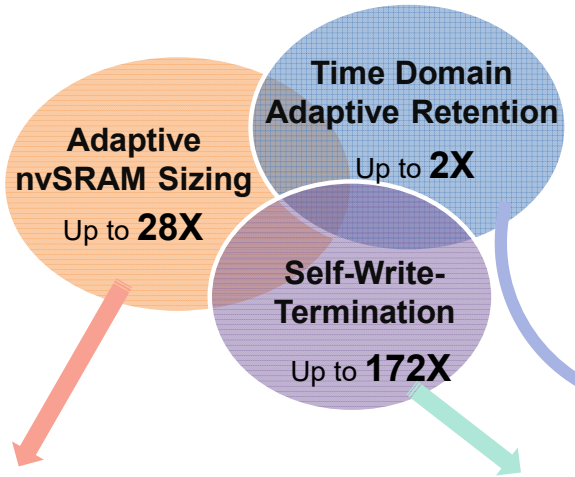

Restore Time/Energy under Different Configurations

Configuration		1422 nvFF	1422 nvFF 16B nvSRAM	1422 nvFF 256B nvSRAM	1422 nvFF 1KB nvSRAM	1422 nvFF 4KB nvSRAM
Restore time/ns	1WL Parallelism	20	20	170	--	--
	4WL Parallelism	20	20	50	170	--
	16WL Parallelism	20	20	20	50	170
Restore Energy/pJ		21.3	27.1	113	390	1496

Store Energy vs. nvSRAM capacity

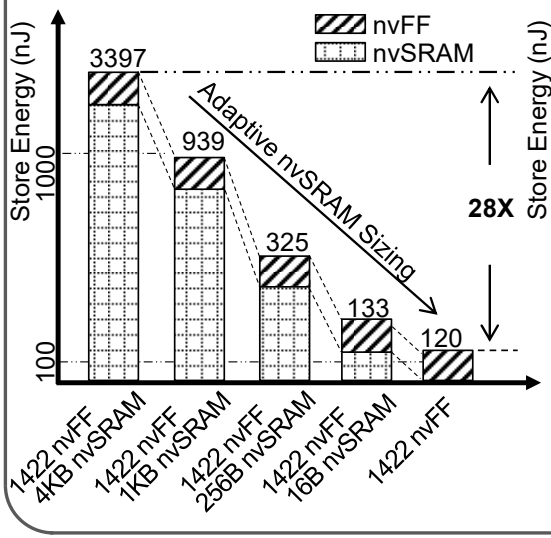

Store Energy vs. Match Ratio

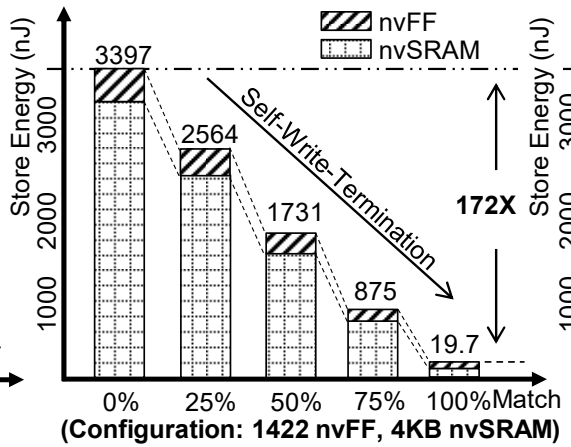

(Configuration: 1422 nvFF, 4KB nvSRAM)

Store Energy vs. Short Power Interrupt Rate *

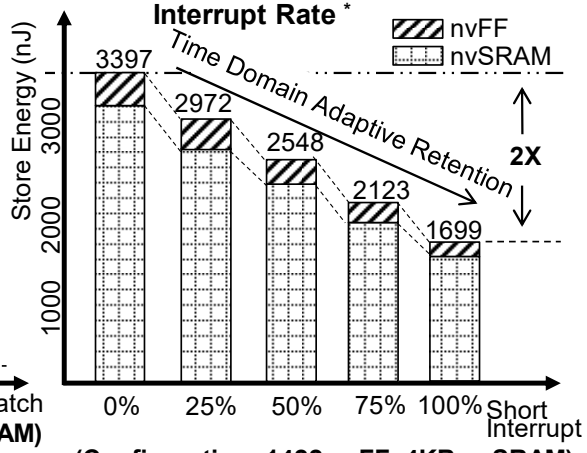

(Configuration: 1422 nvFF, 4KB nvSRAM)

* Average Short Power Interrupt Time = ½ Breakeven Time

Metric Comparison

Metric	ESSCIRC'12	ISSCC'13	ISSCC'14	This Work
Technology (nm)	130	130	90	65
NVP Features	8051	Cortex-M	MSP430	8051
Memory Tech.	FeRAM	FeRAM	MRAM	ReRAM
NVM Features	2-macro, SRAM+FeRAM	1-macro, FeRAM	1-macro, SpinRAM	1-macro nvSRAM
Frequency (MHz)	25	8	20	>100 (4X)
nvFF Restore Time (ns)	3000	384	120	<20 (6X)
Sytem Restore Time	2.95 ms	Store Data in FeRAM; Low Speed & High Power	Store Data in SpinRAM; Low Speed & High Power	<170 ns
System Store Time	2.95 ms			<1.02 ms
System Restore Energy	44562 nJ			0.45 nJ (>6000X)
System Store Energy	44.56 uJ			0.40 uJ

Conclusion

- NVP is promising for energy harvesting IoTs.
- The first ReRAM-based NVP integrates SWT-supported nvSRAM and nvFFs.
- 1-macro hybrid nvSRAM enables 4× faster clock frequency than state-of-the-art NVPs.
- Adaptive data retention and SWT schemes achieve 6× speedup and >6000× energy reductions in restore operations.

Acknowledgement

- THU-NTHU Joint Project
- 863 Project 2013AA01320
- Beijing High-Caliber Talents Project YETP0102
- Beijing High-Tech Innovation Center
- NSFC Grant #61271269